


RESEARCH

Open Access



Online reactions to the 2017 'Unite the right' rally in Charlottesville: measuring polarization in Twitter networks using media followership

Joseph H. Tien^{1*} , Marisa C. Eisenberg², Sarah T. Cherng³ and Mason A. Porter⁴

*Correspondence: tien.20@osu.edu

¹Department of Mathematics and Mathematical Biosciences Institute; The Ohio State University, Columbus OH, USA
Full list of author information is available at the end of the article

Abstract

Network analysis of social media provides an important new lens on politics, communication, and their interactions. This lens is particularly prominent in fast-moving events, such as conversations and action in political rallies and the use of social media by extremist groups to spread their message. We study the Twitter conversation following the August 2017 'Unite the Right' rally in Charlottesville, Virginia, USA using tools from network analysis and data science. We use media followership on Twitter and principal component analysis (PCA) to compute a 'Left'/'Right' media score on a one-dimensional axis to characterize Twitter accounts. We then use these scores, in concert with retweet relationships, to examine the structure of a retweet network of approximately 300,000 accounts that communicated with the #Charlottesville hashtag. The retweet network is sharply polarized, with an assortativity coefficient of 0.8 with respect to the sign of the media PCA score. Community detection using two approaches, a Louvain method and InfoMap, yields communities that tend to be homogeneous in terms of Left/Right node composition. We also examine centrality measures and find that hyperlink-induced topic search (HITS) identifies many more hubs on the Left than on the Right. When comparing tweet content, we find that tweets about 'Trump' were widespread in both the Left and Right, although the accompanying language (i.e., critical on the Left, but supportive on the Right) was unsurprisingly different. Nodes with large degrees in communities on the Left include accounts that are associated with disparate areas, including activism, business, arts and entertainment, media, and politics. By contrast, support of Donald Trump was a common thread among the Right communities, connecting communities with accounts that reference white-supremacist hate symbols, communities with influential personalities in the alt-right, and the largest Right community (which includes the Twitter account **FoxNews**).

Keywords: United States politics, Political extremism, Media polarization, Social media, Twitter, Community structure, Principal component analysis

Introduction

On 11–12 August 2017, a ‘Unite the Right’ rally was held in Charlottesville, Virginia, USA in the context of the removal of Confederate monuments from nearby Emancipation Park. Attendees at the rally included people in the ‘alt-right’, white supremacists, Neo-Nazis, and members of other far-right extremist groups (Fausset and Feuer 2017). At this rally, there were violent clashes between protesters and counter-protesters. A prominent event amidst these clashes was the death of Heather Heyer when a rally attendee rammed his car into a crowd of counter-protesters (Duggan and Jouvenal 2019). In the aftermath, President Donald Trump stated that there were ‘very fine people on both sides’ (Shear and Haberman 2017). White supremacists were galvanized by Trump’s response, with one former leader stating that the president’s comments marked “the most important day in the White nationalist movement” (Daniels (2018), p. 61). Reactions to the removal of Confederate statues, the violence at the rally, and President Trump’s controversial response generated vigorous debate.

In this paper, we present a case study of the structure of the online conversation about Charlottesville in the days following the ‘Unite the Right’ rally. These data allow one to study far-right extremism in the context of broader public opinion. Did support for President Trump’s handling of Charlottesville extend beyond white supremacists and the ‘alt-right’? Did the response to Charlottesville split simply along partisan lines, or was the reaction more nuanced? We examine these questions with tools from network analysis and data science using Twitter data from communication following the ‘Unite the Right’ rally that included the hashtag #Charlottesville. Our specific objectives are to (1) present a simple approach for characterizing Twitter accounts based on their online media preferences; (2) use this characterization to examine the extent of polarization in the Twitter conversation about Charlottesville; (3) evaluate whether key accounts were particularly influential in shaping this discussion; (4) identify natural groupings (in the form of network ‘communities’) of accounts based on their Twitter interactions; and (5) characterize these communities in terms of their account composition and tweet content.

Social media platforms are important mechanisms for shaping public discourse, and data analysis of social media is a large and rapidly growing area of research (Tüfekci 2014). It has been estimated that almost two thirds of American adults use social media platforms (Perrin 2015), with even higher usage among certain subsets of the population (such as activists (Tüfekci 2017) and college students (Perrin 2015)). Online forums and social media platforms are also significant mechanisms for communication, dissemination, and recruitment for various types of ethnonationalist and extremist groups (Daniels 2018). Twitter, in particular, has been a key platform for white-supremacist efforts to shape public discourse on race and immigration (Daniels (2018), p. 64).

As a network, Twitter encompasses numerous types of relationships. It is common to analyze them individually as *retweet networks* (e.g., see Conover et al. (2011); Romero et al. (2011)), *follower networks* (e.g., see Colleoni et al. (2014)), *mention networks* (e.g., see Conover et al. (2011)), and others. There is an extensive literature on Twitter network data, and the myriad topics that have been studied using such data include political protest and social movements (Barberá et al. 2015; Beguerisse-Díaz et al. 2014; Cihon and Yasseri 2016; Freelon et al. 2016; Lynch et al. 2014; Tremayne 2014; Tüfekci and Wilson 2012), epidemiological surveillance and monitoring of health behaviors (Chew and

Eysenbach 2010; Denecke et al. 2013; Lee et al. 2014; McNeill et al. 2016; Salathé et al. 2013; Signorini et al. 2011; Towers et al. 2015), contagion and online content propagation (Lerman and Ghosh 2010; Weng et al. 2013), identification of extremist groups (Benigni et al. 2017), and ideological polarization (Conover et al. 2011; Garimella and Weber 2017; Morales et al. 2015). The combination of significance for public discourse, data accessibility, and amenability to network analysis makes it very appealing to use Twitter data for research. However, important concerns have been raised about biases in Twitter data (Cihon and Yasseri 2016; Mitchell and Hitlin 2014; Morstatter et al. 2014) in general and hashtag sampling in particular (Tüfekci 2014), and it is important to keep them in mind when interpreting the findings of both our study and others.

There are also many studies of how the internet and social media platforms affect public discourse (Bail et al. 2018; Flaxman et al. 2016; Sunstein 2001). In principle, social media and online news consumption have the potential to increase exposure to disparate political views (Wojcieszak and Mutz 2009). In practice, however, they instead often serve as filter bubbles (Allcott and Gentzkow 2017; Pariser 2011) and echo chambers (Flaxman et al. 2016; Sunstein 2001); they thus may heighten polarization. Several previous studies have examined political homophily in Twitter networks (Colleoni et al. 2014; Conover et al. 2011; Feller et al. 2011; Morales et al. 2015; Pennacchiotti and Popescu 2011), and there has been some analysis of tweet content and followership of political accounts (Colleoni et al. 2014). We also examine political polarization using Twitter data, but we take a different approach: we focus on the homophily of media preferences on Twitter. Specifically, we examine media followership on Twitter and perform principal component analysis (PCA) (Jolliffe 2002) to calculate a scalar measure of media preference. We then use this scalar measure to characterize accounts in our Charlottesville Twitter data set. To study homophily, we examine assortativity of this scalar quantity for accounts that are linked by one or more retweets. Our approach is conceptually simple, has minimal data requirements (e.g., there are no training data sets), and is straightforward to implement. The use of media followership is also appealing for political studies, as it is known that media preferences correlate with political affiliation (Budak et al. 2016; Mitchell et al. 2014).

The influence of Twitter accounts on shaping content propagation and online discourse depends on many factors, including the number of 'followers' (accounts who subscribe to a given account's posts, which then appear in their feed), community structure and other aspects of network architecture (Weng et al. 2013), tweet activity (and other account characteristics) (Barberá et al. 2015), and specific tweet content (Cha et al. 2010). One can calculate 'centrality' measures (Newman 2018) to identify important nodes in a Twitter network. There are many notions of centrality, including degree, PageRank (Brin and Page 1998), betweenness (Freeman 1977), and hyperlink-induced text search (HITS, which allows the examination of both hubs and authorities) (Kleinberg 1999). In the context of our study, it is also useful to keep in mind that some structural features are particular to Twitter networks; these may influence which centrality measures are most appropriate to consider. Prominent examples of such features include asymmetry between the numbers of followers and accounts that are followed for many accounts (Cha et al. 2010), automated accounts ('bots') that may retweet at very high frequencies (Davis et al. 2016), and heterogeneous retweeting properties across different accounts (Romero et al. 2011). The importance of such features has also led to the development of Twitter-specific centrality measures (Bougoussa and Ben Romdhane 2015; Romero et al. 2011; Weng et al. 2010).

In our investigation, we examine a variety of different measures of centrality for the #Charlottesville retweet network to identify important accounts both with respect to the generation of novel content and with respect to the spreading of existing content.

Community detection, in which one seeks dense sets (called ‘communities’) of nodes that are connected sparsely to other dense sets of nodes, is another approach that can give insights into network structure (especially at large scales) (Porter et al. 2009; Fortunato and Hric 2016). Communities in a network can influence dynamical processes, such as content propagation (Jeub et al. 2015; Salathé and Jones 2010; Weng et al. 2013). Investigating community structure and other large-scale network structures can be very useful for the study of online social networks, as some accounts are anonymous and demographic data may be incomplete or of questionable validity. Community detection is helpful for discovering tightly-knit groupings of accounts that can help reveal what segments of a population are engaged in a conversation on Twitter. One can then examine such groupings, in conjunction with other tools from network analysis, to characterize communities in terms of structural network properties (e.g., distributions of degree or other centrality measures) and/or metadata (e.g., profile information), identify influential accounts within communities, and study dynamical processes on a network (such as how content propagates both within and between communities (Weng et al. 2013)).

In the present paper, we combine community detection with analysis of tweet content within and across communities. Previous studies have reported that there are differences in language in different online communities (Bryden et al. 2013). Such differences can help reveal differences in demography, political affiliation, and views on specific topics (Colleoni et al. 2014; Conover et al. 2011; Wong et al. 2016). For example, the ‘linguistic framing’ of issues such as immigration can help reveal political orientations and agendas (Huber 2009; Lakoff and Ferguson 2006), and changes in language over time can reflect political movements and influence campaigns (Morgan 2017). We combine community detection with tweet-content analysis to compare subsets of the Twitter population who participated in the #Charlottesville conversation by characterizing them based on the language that was used in different communities for describing both the broader conversation topic (namely, #Charlottesville) and specific subtopics (e.g., ‘Trump’).

Our paper proceeds as follows. In “[Data collection](#)”, we briefly discuss our Twitter data collection and cleaning. In “[Twitter media preferences](#)”, we discuss how we characterize nodes based on media preference and PCA. In particular, we show that the first principal component provides a good classification of nodes as ‘Left’ (specifically, nodes with a negative media-preference score) or ‘Right’ (specifically, nodes with a positive media-preference score) with respect to their media preference. In “[Network structure](#)”, we examine the structure — in terms of both centrality measures and large-scale community structure — of a network of retweet relationships that we construct from our Twitter data. We compare central nodes on the Left and Right, and we also examine Left/Right node composition within communities in the retweet network. In “[Media-preference assortativity](#)”, we examine the media-preference assortativity of nodes in the retweet network. Our results allow us to gauge the extent to which the Twitter conversation, with respect to who retweets whom, splits according to media preferences. In “[Comparison of tweet content between Left and Right](#)”, we examine differences in tweet content between the Left and Right communities. Taken together, our results illustrate the extent of polarization in the Twitter conversation

about #Charlottesville. In “[Conclusions and discussion](#)”, we conclude and discuss our results.

Data collection

We collected Tweets with the hashtag #Charlottesville and the follower lists for 13 media organizations using Twitter’s API (application program interface). Public data accessibility through Twitter’s API has greatly facilitated investigations of Twitter data, but such data have important limitations (Cihon and Yasseri 2016; Tüfekci 2014), including potential biases due to Twitter’s proprietary API sampling scheme (Cihon and Yasseri 2016). For example, Morstatter et al. (2014) illustrated that Twitter’s API can produce artifacts in topical tweet volume, potentially resulting in misleading changes in the number of tweets on a given topic over time. In our investigation, we do not consider changes in tweet volume over time; instead, we examine features of the data after aggregating over a collection-time window.

There are also important considerations in the use of hashtags. Tüfekci (2014) discussed several potential issues with hashtag sampling, including different hashtag usages across different groups and discontinuation of a given hashtag once the corresponding topic has been established. (This latter phenomenon is called ‘hashtag drift’ (Salganik 2017).) We collected the tweets that we study from a six-day period from shortly after the ‘Unite the Right’ rally; this should lessen the potential for hashtag drift. As was pointed out by Tüfekci (2014), hashtag sampling draws from accounts that choose to tweet a given hashtag, and this necessarily entails biases. Nevertheless, hashtag sampling is able to provide valuable insights into the shape of online conversations. For example, we use the collected data to examine what types of accounts chose to post tweets about #Charlottesville. It is known, for example, that the extent that ‘peripheral’ accounts engage in online conversations about social protests can be an important factor for content propagation on Twitter (Barberá et al. 2015).

Our data collection is in accord with the Twitter Terms of Service and Developer Agreement. To protect user privacy, we include account names (i.e., “handles”) only for Twitter-verified accounts and Twitter accounts that belong to organizations. As described by Twitter, “an account may be verified if it is determined to be an account of public interest” (Twitter, Inc. 2019). We have posted network data (without account names or tweet content), together with code for analyzing the structure of these data, at <https://osf.io/487fw/>.

Tweets about #Charlottesville

We used Twitter’s search API to sample 486,894 publicly available tweets that include the hashtag #Charlottesville and were posted by 270,975 unique accounts between 16 August 2017 and 21 August 2017. Our data includes account name (i.e., “handle”), time and date in coordinated universal time (UTC), and tweet content. In UTC, the earliest tweet date is 2017-08-16 22:16:21, and the latest tweet date is 2017-08-20 01:48:00. We performed our data acquisition using the Python package `tweepy`.

Media followership

In December 2016, we used the Twitter API to acquire the complete lists of Twitter users who follow the following 13 media accounts: `BreitbartNews`, `DRUDGE_REPORT`,

Table 1 First principal component of the media-followership matrix M

Media account	Value of the first principal component
BreitbartNews	0.4071
DRUDGE_REPORT	0.3843
FoxNews	0.3779
theblaze	0.2054
NRO	0.0970
csmonitor	-0.0235
WSJ	-0.1183
FiveThirtyEight	-0.1520
dailykos	-0.1893
thenation	-0.2362
MotherJones	-0.3115
washingtonpost	-0.3321
NPR	-0.3945

The red entries designate a positive entry. For each node, we use its value in the first principal component to characterize it in the retweet network: 'Left' refers to nodes with a negative first-principal-component value, and 'Right' refers to nodes with a positive first-principal-component value.

FiveThirtyEight, FoxNews, MotherJones, NPR, NRO¹, WSJ², csmonitor, dailykos, theblaze, thenation, and washingtonpost. At the time of access, these media accounts had significant Twitter followings, ranging from 62,078 followers (csmonitor) to more than 12 million followers (WSJ). They include both sources that studies have concluded as preferred by conservative readers and sources that they have concluded as preferred by liberal ones (Budak et al. 2016; Mitchell et al. 2014).

Twitter media preferences

In this section, we use media preferences on Twitter to characterize nodes in our #Charlottesville data set. Specifically, we find that using PCA provides an effective characterization of nodes as 'Left' or 'Right'. In subsequent sections, we use this characterization to examine media-preference polarization in the Twitter conversation about Charlottesville.

Of the Twitter accounts in our #Charlottesville data set, 99,412 accounts followed at least one of the 13 media sources that we listed in "Media followership" at the time (December 2016) that we accessed the media follower lists. Restricting to these accounts gives a $99,412 \times 13$ media-choice matrix M of 0 entries (not following) and 1 entries (following). We perform PCA on M , and we give the first component in Table 1. We use the 'standard' type of PCA in our investigation (Jolliffe 2002). For a discussion of a variant of PCA that is designed for Boolean data, see Landgraf and Lee (2015).

We interpret the first component as encoding liberal versus conservative media preference, as reflected by the signs of the entries of this component. Specifically, media accounts with a positive value in the first principal component (PC) seem to correspond to accounts that previous studies have found to have a conservative slant (and to be preferred by individuals who identify as conservative), whereas accounts with a negative value in the first PC correspond predominantly to accounts that studies have concluded to have a liberal slant and/or are preferred by liberals (Budak et al. 2016; Mitchell et al. 2014). The sign of the score in the first PC is also consistent with conventional

¹NRO is the Twitter account for *The National Review*.

²WSJ is the Twitter account for *The Wall Street Journal*.

wisdom about liberal versus conservative leanings of these media accounts, with the exception of *The Wall Street Journal* (WSJ), which is widely considered to be conservative-leaning (Flaxman et al. 2016) but has a negative first PC value. However, our findings are consistent with previous studies that, based on readership and co-citations, grouped *The Wall Street Journal* with liberal media organizations (Flaxman et al. 2016; Gentzkow and Shapiro 2011; Groseclose and Milyo 2005). By contrast, previous research that examined article content identified *The Wall Street Journal* as politically conservative (Budak et al. 2016). Although the sign of the first PC value has a clear interpretation, the magnitude of these entries does not appear to provide an intuitive ordering (for example, with respect to a hand-curated media-bias chart (Otero 2018)) on the liberal–conservative spectrum.

In the rest of our paper, we focus on the value of the first PC; for simplicity, we use the term ‘media PCA score’ to refer to this score. Positive values for this score indicate followership of the media accounts that we show in red, whereas negative values indicate followership of accounts that we show in black (see Table 1). To frame our discussion, we refer to nodes with a positive media PCA score as nodes on the ‘Right’ and to those with a negative media PCA score as ones that are on the ‘Left’, although we note that we have not validated this measure as an indicator of political belief or affiliation. Our approach is similar to that of Bail et al. (2018), who applied PCA to a followership network of a large set of ‘opinion leaders’ to assess political orientation.

Network structure

In this section, we explore structural features of our *retweet network* \tilde{G} , which is a weighted, directed graph with weighted adjacency matrix \tilde{A} , where \tilde{A}_{ij} denotes the number of times that node j retweeted node i . In particular, we examine degree distributions (see “Degree distributions”), calculate and compare several different centrality measures (see “Centralities”), and detect communities using two widely-used algorithms (see “Community structure”). We combine these structural features with node characterization according to media preference (see “Media followership”) to (1) examine how central nodes differ between Left and Right and (2) describe communities based on their Left/Right node composition.

The graph \tilde{G} has 238,892 nodes, 365,589 edges, and 389,736 retweets. In the remainder of our paper, we study G , the largest connected component of \tilde{G} when we ignore directionality (so it is \tilde{G} ’s largest weakly connected component). The graph G has 221,137 nodes, 353,548 edges, and 376,978 retweets. Let A denote the weighted adjacency matrix that is associated with G . In all cases, weights represent multi-edges, where the multi-edge from node j to node i corresponds to the number of retweets by account j of any tweet by account i .

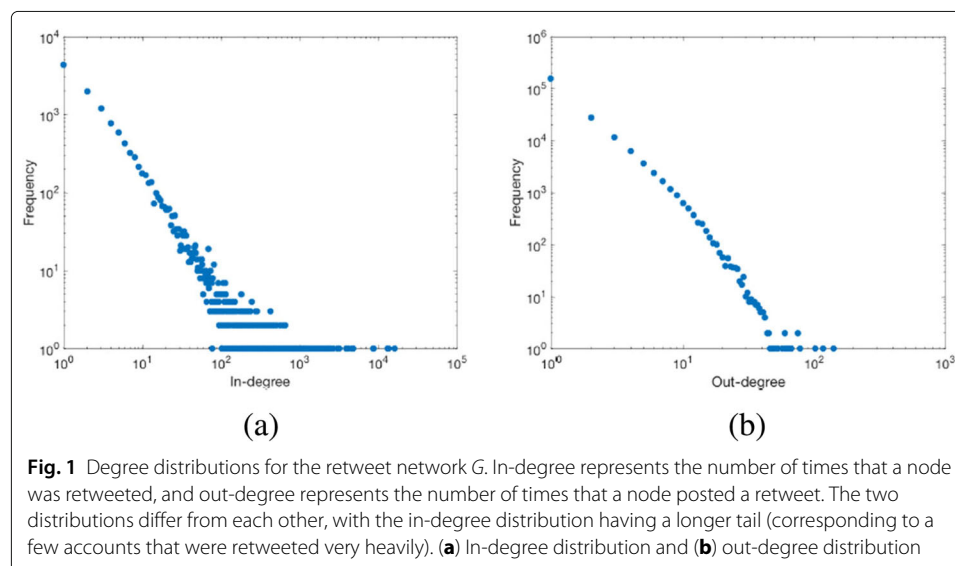
Degree distributions

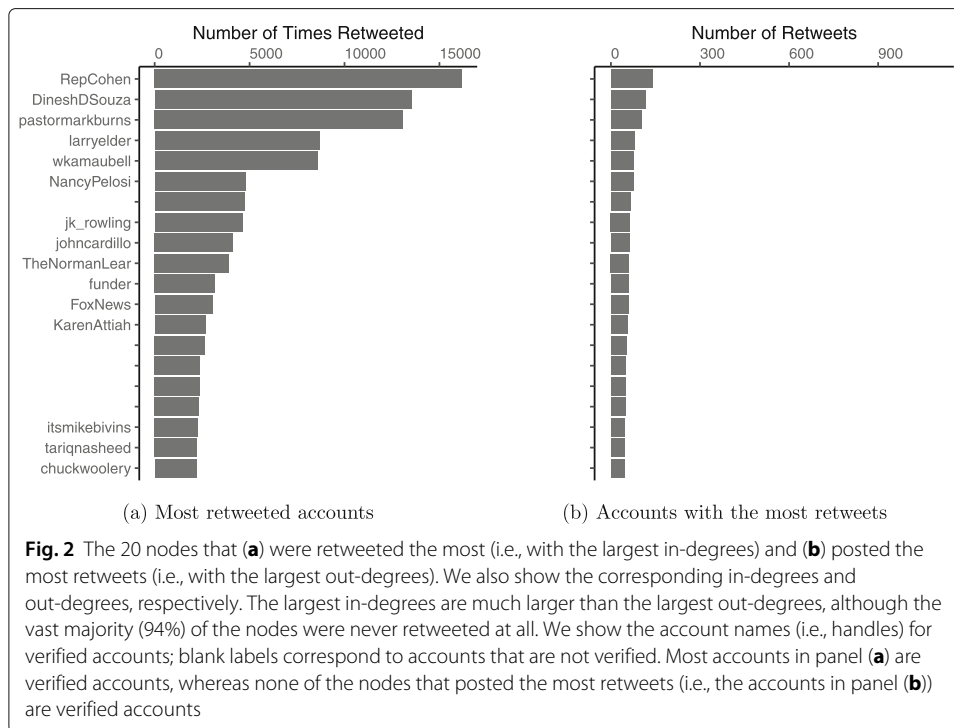
The out-degree of node k corresponds to the total number of retweets that were posted by node k , and the in-degree of k corresponds to the total number of times that node k was retweeted. Unless we specifically note otherwise, we include weights when calculating the in-degrees and out-degrees (i.e., we count all edges in a multi-edge), so we are calculating what are sometimes called “in-strengths” and “out-strengths”. For example, $\sum_{i=1}^n A_{ij}$ gives the out-degree of node j , and $\sum_{j=1}^n A_{ij}$ gives the in-degree of node i . In Fig. 1, we show

the in-degree and out-degree distributions for G . The two distributions differ markedly, as the in-degree distribution has a much longer tail (corresponding to a few accounts that were retweeted very heavily).

In Fig. 2a, we show the in-degrees for the 20 most heavily retweeted accounts. The mean in-degree is 1.70, and the standard deviation is 69.22, indicating extreme heterogeneity in the number of times retweeted. The account (`RepCohen`) with the largest in-degree was retweeted 16,180 times. By contrast, 208,241 nodes (i.e., 94% of them) in G were never retweeted at all. We also observe heterogeneity in the out-degree, but it is much less extreme than for in-degree; the standard deviation is 4.89. (By definition, the mean in-degree and mean out-degree are the same, as every edge has both an origin and terminus in G .) The account with the largest out-degree sent 141 retweets in our data set. By contrast, 7,852 accounts had an out-degree of 0; these accounts were retweeted, but they did not retweet any accounts. In Fig. 2b, we show the 20 accounts that sent the most retweets.

We compare the in-degree and out-degree distributions of accounts with and without media PCA scores to examine whether there are systematic differences between the two types of accounts. (The former are the 99,412 accounts that followed at least one of the 13 focal media sources.) The heterogeneity in the in-degree distribution that we observed when examining all nodes in G is also present when we consider the in-degree distribution separately for nodes with and without media PCA scores; the standard deviation is 105.24 for nodes with a media PCA score and 36.65 for nodes without one. The mean in-degree of nodes with a media PCA score is larger than that of nodes without one (2.85 versus 1.08). Nodes with large in-degree with media PCA scores include `DineshDSouza`, `pastormarkburns`, `RepCohen`, `wkamaubell`, and `johncardillo`. However, there are also some heavily retweeted nodes — such as `larryelder`, `TheNormanLear`, and `NancyPelosi` — that do not follow any of the 13 media accounts that we used for computing media PCA scores. We thus cannot compute media PCA scores for these nodes.



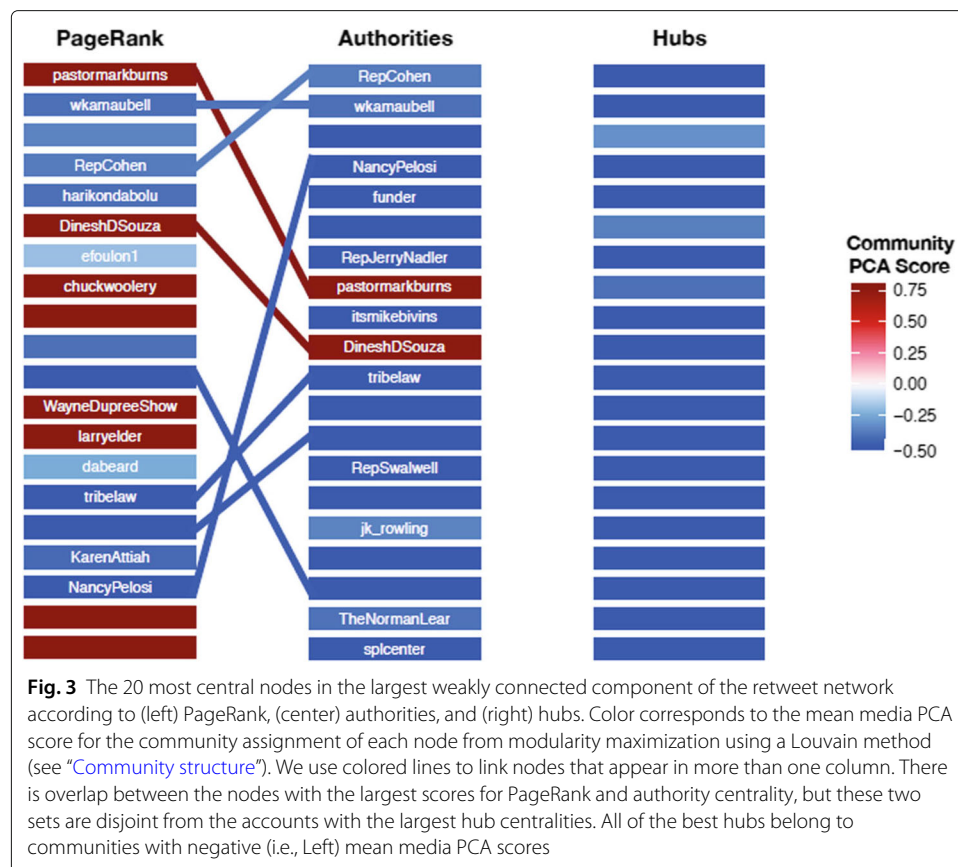


Centralities

We now examine important accounts by computing several centrality measures (Newman 2018; Baek et al. 2019). We start with degree (i.e., degree centrality), the simplest way of trying to measure a node's importance. In Fig. 2, we show the 20 nodes with the largest in-degrees and the 20 nodes with the largest out-degrees. These two sets are disjoint, indicating that the nodes that generated most of the original content in the Twitter conversation about #Charlottesville were distinct from those that were most active in promoting existing content through retweets. Degree is a local centrality measure that does not take into account any characteristics of neighboring nodes. As a comparison, we also calculate two other widely-used centrality measures, PageRank (Brin and Page 1998) and HITS (Kleinberg 1999), that take some non-local information into account.

One obtains PageRank scores from the stationary distribution of a random walk on a network that combines transitions according to network structure and 'teleportation' according to a user-supplied distribution (Gleich 2015), with a parameter that determines the relative weightings of these two processes. We compute PageRank with standard uniform-at-random teleportation using MATLAB's `centrality` function with the default "damping factor" of 0.85 (so teleportation occurs for 15% of the steps in the associated random walk). In the left column of Fig. 3, we list the 20 most central nodes according to PageRank. Nine of these nodes are also on our list of nodes with the largest in-degrees. An exception is `harikondobalu`, which was retweeted only 38 times in our data set. The large PageRank value for `harikondobalu`, despite its small in-degree, reflects the fact that `harikondobalu` was one of only two nodes that were retweeted by `wkamaubell`, which was retweeted 8,582 times.

Hub and authority centralities (Kleinberg 1999) are another useful pair of centrality measures. Using the HITS algorithm, one can simultaneously examine *hubs* and



authorities. As discussed in Kleinberg (1999), a good hub tends to point to good authorities, and a good authority tends to have good hubs that point to it. In the context of retweeting, we expect that accounts with large authority scores tend to be retweeted by accounts with large hub scores, and we expect that good hub accounts tend to retweet accounts that are good authorities. As in PageRank, the importances of adjacent nodes influence a node’s hub and authority scores. With our convention that entry (i, j) of a graph’s adjacency matrix corresponds to the edge weight from node j to node i , hub and authority scores correspond, respectively, to the principal right eigenvectors of $A^t A$ and AA^t . We compute hubs and authorities using MATLAB’s `centrality` function.

We list the 20 nodes with the largest authority and hub scores, respectively, in the center and right columns of Fig. 3. The color of each account indicates the mean media PCA score of its associated community from modularity maximization using a Louvain method (Blondel et al. 2008; Newman and Girvan 2004; Newman 2006) (see “Community structure”). Only two of the nodes among the top-20 authorities are in communities with positive (i.e., Right) mean media PCA scores; these accounts, *pastormarkburns* and *DineshDSouza*, belong to prominent conservative personalities. Neither *pastormarkburns* nor *DineshDSouza* were ever retweeted by any of the top-50 hubs. By contrast, all of the other authorities were retweeted at least three times by the leading hubs. When we consider all nodes, we observe that the hub scores have a bimodal distribution, with a clear separation between the nodes with small and large values (e.g., using 4×10^{-5} as a threshold hub score to separate ‘small’ and ‘large’ values). We refer to nodes with hub scores that are larger than 4×10^{-5} as ‘large

hub-score nodes'. For example, among the set of nodes that retweeted *DineshDSouza*, the fraction that are large hub-score nodes is 9.0×10^{-4} . The fraction of nodes that retweeted *pastormarkburns* that are large hub-score nodes is 1.0×10^{-3} . The fraction of nodes that retweeted *itsmikebivins* (an account in a Left-leaning community whose authority score lies between those of *pastormarkburns* and *DineshDSouza*) that are large hub-score nodes is 1.3×10^{-2} . A few other examples of such fractions are 0.05 for nodes that retweeted *wkamaubell*, 0.15 for nodes that retweeted *tribelaw*, and 1 for nodes that retweeted *RepCohen*.

There are two qualitatively different ways for a node to have a large authority score: it can either be retweeted many times (e.g., *DineshDSouza*), or it can be retweeted by nodes with large hub scores (e.g., *itsmikebivins*). Both of the large-authority Right accounts (*DineshDSouza* and *pastormarkburns*) lie in the former category.

Using Fig. 3, we briefly compare important accounts according to different centrality measures. As one can see in Fig. 3, there is some overlap between the top-PageRank and top-authority accounts. Note, however, that fewer than half of the top-PageRank accounts are also among the top-authority accounts. By contrast, the set of top hubs is disjoint from the top-PageRank and top-authority accounts in Fig. 3. Additionally, more than half of the top-PageRank and top-authority accounts in Fig. 3 are verified accounts, whereas none of the top-hub accounts are verified accounts.

Community structure

To examine large-scale structure in the #Charlottesville retweet network, we use community detection to identify tightly-knit sets (so-called 'communities') of accounts with relatively sparse connections between these sets (Porter et al. 2009; Fortunato and Hric 2016). In our investigation, we employ two widely-used community-detection methods: modularity maximization (Newman and Girvan 2004; Newman 2006) and InfoMap (Rosvall and Bergstrom 2008). A major challenge in community detection is parsing what results reflect a network's features, rather than artifacts from a community-detection method. Modularity maximization and InfoMap are two methods, which use rather different approaches from each other, that have been used successfully on a wide variety of problems. We expect that broad structural features in a network that we observe using both of these methods are likely to be robust to the particular choice of community-detection method, so we expect them to be actual features of the data (rather than artifacts). There exist many other community-detection methods, including statistical inference via stochastic block models (Peixoto 2017) and local methods based on personalized PageRank (Jeub et al. 2015; Gleich 2015). Exploring our retweet network with other community-detection methods is outside the scope of the present article, but we encourage readers to explore our data set with them. It is available at <https://osf.io/487fw/>.

Modularity maximization

The modularity of a particular assignment of a network's nodes into communities measures the amount of intra-community edge weight, relative to what one would expect at random under some null model (Newman and Girvan 2004; Newman 2006). Modularity maximization treats community detection as an optimization problem by seeking an assignment of nodes to communities that maximizes a modularity objective function. A version of modularity for weighted, directed graphs is (Arenas et al. 2007; Leicht and Newman 2008)

$$Q = \frac{1}{w} \sum_{i=1}^n \sum_{j=1}^n \left(A_{ij} - \gamma \frac{w_i^{\text{in}} w_j^{\text{out}}}{w} \right) \delta(C_i, C_j), \quad (1)$$

where

$$w = \sum_{i=1}^n \sum_{j=1}^n A_{ij} \quad (2)$$

is the sum of all edge weights in a network; w_k^{in} and w_k^{out} are the in-strength (i.e., a weighted in-degree) and out-strength (i.e., weighted out-degree), respectively, of node k ; the community assignment of node k is C_k ; the symbol δ denotes the Kronecker delta; and γ is a resolution parameter that controls the relative weight of the null model (Lambiotte et al. 2015). Our null-model matrix elements are $P_{ij} = \frac{w_i^{\text{in}} w_j^{\text{out}}}{w}$, so this null model is a type of configuration model (Fosdick et al. 2018). In this null model, we preserve expected in-strength and expected out-strength, but we otherwise randomize connections (Porter et al. 2009). For most of our computations, we use the resolution-parameter value $\gamma = 1$ as a default. However, in “Community characteristics across different resolution-parameter values”, we compare results using a few values of γ (that span three orders of magnitude).

To maximize Q , we use a GENLOUVAIN variant (Jeub et al. 2011–2016) (which is implemented in MATLAB and was released originally in conjunction with Mucha et al. (2010)) of the locally-greedy Louvain algorithm (Blondel et al. 2008). To use the code from Jeub et al. (2011–2016), we symmetrize the modularity matrix B , where $B_{ij} = A_{ij} - \gamma \frac{w_i^{\text{in}} w_j^{\text{out}}}{w}$. As discussed in Leicht and Newman (2008), this is distinct from symmetrizing the adjacency matrix A .

Modularity maximization using GENLOUVAIN yields 228 communities, which range in size from 2 nodes to 47,321 nodes.

InfoMap

InfoMap is a community-detection method that is based on the flow of random walkers on graphs (Rosvall and Bergstrom 2008).³ It uses the intuition that a random walker tends to be trapped for long periods of time within tightly-knit sets of nodes (Fortunato and Hric 2016). Rosvall and Bergstrom (2008) made this idea concrete by trying to minimize the expected description length of a random walk. For example, one can obtain a concise description of a random walk by allowing node names to be reused between communities. One can apply InfoMap to weighted, directed graphs; and it has been used previously to study Twitter data (Weng et al. 2013). To study a directed graph, one introduces a teleportation parameter (as in PageRank); we use the default teleportation value of $\tau = 0.15$ (Rosvall and Bergstrom 2008).

Our implementation uses code from Edler and Rosvall (2019). With InfoMap, we find 205 communities, which range in size from 1 node to 122,504 nodes.

Large-scale structure of the retweet network

Several features are evident in our community-detection results from both modularity maximization and InfoMap: (1) communities are largely segregated by media PCA score; (2) overall, the communities skew towards the Left; and (3) most of the nodes on the Right are assigned to a large community that includes prominent right-wing personalities and FoxNews.

³One can also interpret modularity maximization in terms of random walks on graphs (Lambiotte et al. 2015).


To examine the relationship between community structure and Left/Right media preference, we compute the mean media PCA score of each community. The proportion of communities with at least one node with a media PCA score is very similar for modularity maximization (204/228; 89%) and InfoMap (183/205; 89%). We also examine the extent of overlap of Left and Right accounts within communities by computing the Shannon diversity index H^s (Shannon 1948) for each community s . This index is given by

$$H^s = - \sum_{i=1}^2 p_i^s \ln p_i^s, \quad (3)$$

where p_1^s and p_2^s (with $p_1^s + p_2^s = 1$ for each s) are the fractions of accounts in community s with Left and Right media preferences, respectively. In Fig. 4, we show the Shannon diversity scores versus mean media PCA scores for the communities that we detect using modularity maximization and InfoMap.

Both community-detection methods yield a predominantly unimodal shape for the relation between PCA score diversity and mean media PCA score, with more extreme mean media PCA scores associated with lower diversity within a community. Communities with ‘centrist’ mean media PCA scores (i.e., ones that are near 0) have relatively small sizes. By contrast, the largest communities tend to have mean media scores that are farther from 0; they also have small Shannon diversity. For example, InfoMap gives two communities that are much larger than the others. One is on the Left (with 122,504 nodes and a mean media PCA score of -0.43), and the other is on the Right (with 58,185 nodes and a mean media PCA score of 0.74). In these two communities, 91% of the nodes in the largest Left community have negative media PCA scores, compared with 6% in the largest Right community. Similarly, the large communities that we obtain from modularity maximization also have little Left/Right node diversity within communities.

Another prominent feature is that both community-detection approaches yield one community on the Right that is much larger than other communities that have a positive mean media PCA score. Additionally, both methods yield a similar set of large-degree accounts in the largest Right community. Specifically, the five nodes with largest in-



degrees and out-degrees are the same, with DineshDSouza, pastormarkburns, larryelder, johncardillo, and FoxNews as the five most heavily retweeted accounts (i.e., the ones with the largest in-degrees) in this community.

Our results in Fig. 4 also suggest that there are more Left-leaning communities than Right-leaning ones. For example, 106/130 (i.e., about 82%) of the InfoMap communities with at least ten nodes have negative mean media PCA scores. Modularity maximization gives a bimodal distribution of community sizes. We refer to communities with at most 100 nodes as ‘small’, communities with more than 100 and at most 1000 nodes as ‘medium-sized’, and communities with more than 1000 nodes as ‘large’. Of the medium-sized and large modularity-maximization communities, 76/93 (i.e., 82%) have negative mean media PCA scores. To give some context, we have PCA scores for 78,339 nodes, and 44,797 of them (about 57%) have a negative media PCA score.

Finer features of the retweet network

One notable difference between the two community-detection methods is that two large communities dominate for InfoMap (one each on the Left and Right), whereas modularity maximization yields a partition of the retweet network that includes many more large communities. We now examine some of the finer details in the large communities that we obtain from modularity maximization.

Modularity maximization yields 41 communities with at least 1,001 nodes. To further characterize these 41 communities, we examine the accounts with the largest in-degrees (i.e., the ones that are retweeted the most) within each community and characterize these nodes by hand from their profiles and, when available (e.g., when account owners are known public personalities), information about the owners of these accounts. More than 85% (specifically, 35 of 41) of these communities have negative (i.e., Left-leaning) mean media PCA scores. The accounts with the largest in-degrees in these 35 communities include activists (e.g., Everytown, IndivisibleTeam, UNHumanRights, and womensmarch), businesses (e.g., benandjerrys), people from arts and entertainment (e.g., jk_rowling, LatuffCartoons, FallonTonight, ladygaga, Sethrogen, TheNormanLear, and wkamaubell), journalists (e.g., AmyKNelson), media organizations (e.g., AJEnglish, CBSThisMorning, and HuffPostCanada), and politicians (e.g., NancyPelosi, RepCohen, and JoeBiden). By contrast, only six of the largest communities have positive (i.e., Right-leaning) mean media PCA scores. The largest of these (with 47,321 nodes) includes opinion leaders on the Right (e.g., DineshDSouza, pastormarkburns, and larryelder) and FoxNews, as we discussed previously. Another community has a mean PCA score close to 0 (specifically, it is 0.086), and it appears to be a business-oriented community with tweets that are critical of Donald Trump. Two of the remaining four communities with positive media scores are Right-oriented activist communities. One activist community has 3,987 nodes, and one of its accounts with among the largest in-degrees (i.e., that is retweeted very heavily) references an influential alt-right account (O’Brien 2018). The other activist community has 2,710 nodes, and one of its most retweeted accounts references a well-known white-supremacist hate symbol in its handle. A third community appears to be a media community with foreign media personalities (e.g., KTHopkins), and the final community of these four is a community that is dominated by accounts that tweet in German.

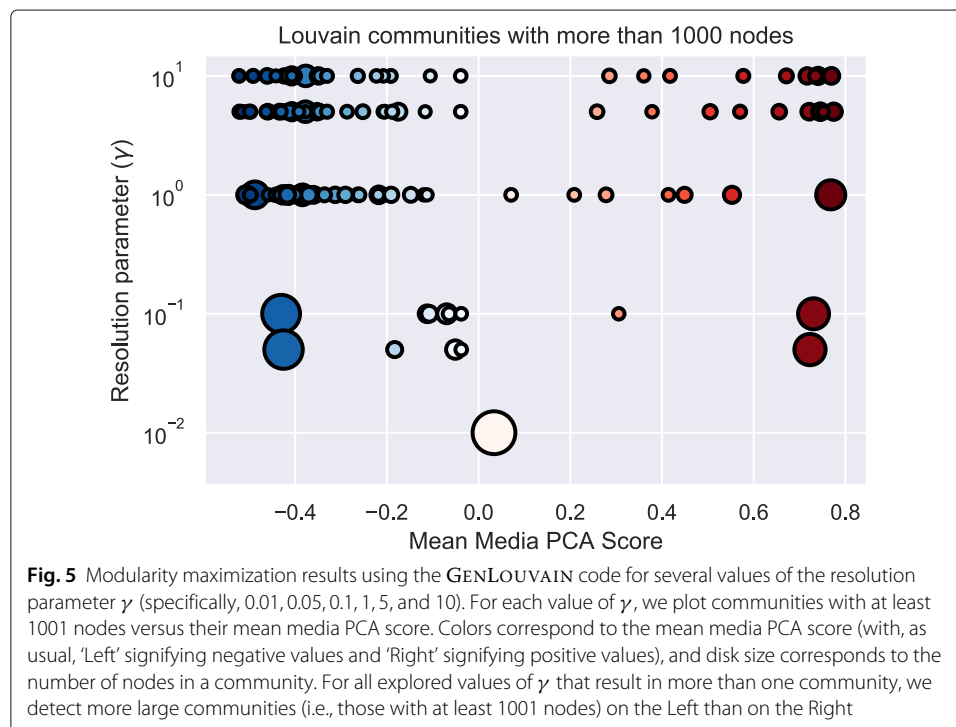
Community characteristics across different resolution-parameter values

To examine the robustness of our findings about community structure in the retweet network, we also conduct modularity maximization using the GENLOUVAIN code for several values of the resolution parameter γ in (1). There is a ‘resolution limit’ for the smallest detectable community size when using modularity maximization, and the size scales of communities that result from modularity maximization can also influence the sizes of the largest communities (Good et al. 2010; Fortunato and Barthélemy 2007). Biases in the sizes of detected communities can skew interpretation of the results of community detection, and distinguishing which results reflect features of a network and which arise from a specific approach or algorithm for community detection is a major challenge. We partially address these concerns by identifying features that are robust to the choice of resolution-parameter value γ .

In Fig. 5, we show our results from community detection using the GENLOUVAIN code with resolution-parameter values γ that range from 10^{-2} to 10. Smaller values of γ in (1) favor fewer communities (with 783 communities using $\gamma = 10$, compared with only a single community for $\gamma = 10^{-2}$). For each value of γ , we plot communities with at least 1001 nodes versus their mean media PCA score. We observe that there are more large (i.e., with at least 1001 nodes) Left-leaning than Right-leaning communities for all examined values of γ for which we detected more than one community, suggesting that this is a robust feature of our Twitter retweet network.

Modular centrality

Our results in “Community structure” suggest that the Charlottesville retweet network has meaningful modules. Ghalmane et al. (2019) examined a way to exploit community structure when identifying influential nodes, and they computed different ‘modular centrality’ measures to describe influence within versus between communities. (See, e.g.,



(Guimerà and Amaral 2005; Fenn et al. 2012) for other work on quantifying centrality values within versus between communities.) In this section, we compute a slightly modified version of the ‘modular degree centrality’ from (Ghalmame et al. 2019). We separately count the number of times (producing an ‘inter-community in-degree’) that a given node is retweeted by nodes that belong to communities other than the one that includes that node and the number of retweets (producing an ‘intra-community in-degree’) by nodes that belong to the same community as that node. We then compute the ratio of the inter-community in-degree to the intra-community in-degree.

In Table 2, we list the ratio of the inter-community in-degree to intra-community in-degree for the 20 nodes with the largest in-degrees in our retweet network. This ratio tends to be significantly larger for nodes that belong to Left-leaning communities (their range is 0.18–0.43, with a mean of 0.27) than for nodes that belong to Right-leaning communities (their range is 0.036–0.078, with a mean of 0.051).

Media-preference assortativity

In this section, we quantify the extent to which retweets occur between nodes with similar media preferences. We do this by examining assortativity in our retweet network in terms of both the value and the sign of the media PCA score. Large assortativity indicates that the Charlottesville conversation largely splits according to media PCA score. Combined with our results in “[Comparison of tweet content between Left and Right](#)” (in which we

Table 2 Number of inter-community retweets divided by the number of intra-community retweets for the 20 most heavily retweeted nodes in the retweet network

Node	Inter-community in-degree divided by intra-community in-degree
RepCohen	0.36725
DineshDSouza	0.040400
pastormarkburns	0.036085
larryelder	0.043745
wkamaubell	0.42938
NancyPelosi	0.21463
	0.25073
jk_rowling	0.22067
johncardillo	0.078068
TheNormanLear	0.31365
funder	0.17688
FoxNews	0.077005
KarenAttiah	0.32803
	0.28350
	0.23508
	0.050066
	0.043146
itsmikebivins	0.23335
tariqnasheed	0.23772
chuckwoolery	0.041193

We obtain communities by maximizing modularity using the GenLouvain code with a resolution-parameter value of $\gamma = 1$. The color of each node indicates the mean media PCA score of its GenLouvain community. We use blue to signify communities that are on the Left (i.e., ones with a negative mean media PCA score) and red to signify communities that are on the Right (i.e., ones with a positive mean media PCA score). Nodes with large in-degrees on the Left have larger values for the number of inter-community retweets divided by the number of intra-community retweets than is the case for nodes with large in-degrees on the Right. We label the public Twitter accounts.

compare tweet content on the Left and Right), this gives a way to assess the amount of polarization in the Twitter conversation.

To examine homophily in media-preference scores in the Twitter conversation about #Charlottesville, we measure media-preference assortativity by computing the Pearson correlation coefficient of the media PCA score for nodes in the retweet network. Specifically, we compute the correlation of the media PCA score for dyads (i.e., nodes that are adjacent to each other via an edge) in the retweet network. We ignore edge weights, and we restrict our calculations to dyads for which we have a PCA score for both nodes. There are 93,521 such pairs.

The correlation coefficient of the media PCA scores is $\rho \approx 0.67$. For comparison, as a null model, we compute the correlation-coefficient distribution for 100,000 random permutations of the PCA scores of the nodes. Specifically, in each realization, we fix the network and assign the PCA scores uniformly at random to the nodes for which PCA scores were available originally. (For an alternative approach for examining assortativity, see Park and Barabási (2007).) The resulting distribution for the correlation coefficient ρ appears to be approximately Gaussian, with a mean of -1.29×10^{-5} and a standard deviation of 0.0033. The z-score for the measured correlation coefficient of 0.67 is larger than 203, indicating that the retweet network has a statistically significant media-preference assortativity.

We also compute the assortativity coefficient r that was introduced by Newman (2002; 2003). Suppose that there are g types of nodes in the network. Following Newman (2003), we calculate

$$r = \frac{\sum_{\ell=1}^g E_{\ell\ell} - \sum_{\ell=1}^g a_{\ell} b_{\ell}}{1 - \sum_{\ell=1}^g a_{\ell} b_{\ell}}, \quad (4)$$

where $E_{\ell s}$ is the fraction of the edges in a network that emanate from a node of type ℓ and terminate at a node of type s , the quantity $a_{\ell} = \sum_{s=1}^g E_{\ell s}$ is the fraction of the edges that emanate from a node of type ℓ , and $b_s = \sum_{\ell=1}^g E_{\ell s}$ is the fraction of edges that terminate at a node of type s .

To calculate (4) for our retweet network, we classify nodes according to the sign of their media PCA score. In the largest weakly connected component of the retweet network, we have PCA scores for 78,339 nodes, of which 44,797 (i.e., about 57% of them) have a negative media PCA score. The resulting assortativity coefficient is $r \approx 0.80$. We show the mixing matrix E in Table 3. As a comparison, Newman (2003) calculated an assortativity coefficient of 0.62 by ethnicity for the sexual-partner network that was described in Catania et al. (1992).

Five⁴ of the media accounts that we used to compute the media PCA score also appear as nodes in the retweet network G . Of these, FoxNews was retweeted 3049 times, NPR was retweeted 69 times, MotherJones was retweeted 15 times, and csmonitor was retweeted 6 times. Removing these media accounts from G has a negligible effect on the assortativity coefficient r .

Although the assortativity by media PCA score in G is rather large, there are some prominent exceptions. For example, RepCurbelo and SenatorTimScott⁵, the accounts for two Republican members of Congress, were heavily retweeted in

⁴These accounts are csmonitor, MotherJones, theblaze, FoxNews, and NPR.

⁵These are the Twitter accounts for Representative Carlos Curbelo (FL, Republican) and Senator Tim Scott (SC, Republican).

Table 3 Mixing matrix E of the edges between the types of accounts. We characterize the type of an account by the sign of its media PCA score

<i>retweeter</i>	<i>originator</i>	
	Left	Right
Left	0.43	0.057
Right	0.044	0.47

As usual, 'Left' indicates a negative media PCA score, and 'Right' indicates a positive media PCA score. (No nodes have a media PCA score of exactly 0.) Accounts tend to mix with (specifically, be adjacent to) accounts with a media PCA score of the same sign, as indicated by the larger values on the diagonal of the mixing matrix.

Left-leaning communities that we detected using modularity maximization. However, both `RepCurbelo` (0.49) and `SenatorTimScott` (0.12) have positive (i.e., Right) media PCA scores, which is consistent with their affiliation with the Republican party. `RepCurbelo` was the fourth-most retweeted account in a community from modularity maximization with a negative (i.e., Left) mean media PCA score (-0.32). `RepCurbelo`, who spoke out strongly against the events in Charlottesville (Curbelo 2017), was retweeted by 22 accounts. We have PCA scores for nine of these accounts, of which four have media PCA scores on the Left. Similarly, `SenatorTimScott` was the second-most retweeted account in a Left-leaning community (with a mean media PCA score of -0.26) that we obtained from modularity maximization. `SenatorTimScott` was retweeted by 78 accounts, and nearly half (specifically, 20 of 43) of the accounts that retweeted `SenatorTimScott` for which we have PCA scores have negative media PCA scores. We identified `RepCurbelo` and `SenatorTimScott` as accounts that warrant examination by first compiling the list of nodes that were retweeted by accounts with media PCA scores of the opposite sign and then examining this list for prominent accounts. One can further develop this approach (for example, to identify negative or mocking retweets (Tüfekci 2014)), and it may be useful in other situations for identifying accounts that generate communication across ideological or other divides.

Comparison of tweet content between Left and Right

In this section, we examine tweet content from Left-leaning and Right-leaning communities. Comparing word and hashtag frequency allows us to see some of the ways in which the Twitter conversation about Charlottesville differed between the Left and the Right.

We use the Python library NLTK 3.3 to tokenize tweets into words and punctuation. In Table 4, we show the 25 most common words in our data set. We separately consider accounts with negative (i.e., Left) and positive (i.e., Right) media PCA scores after removing stop words.⁶ We do not stem the words in our data set, and we treat different capitalizations as different words. We find some overlap between the Left and Right data sets. For example, tweets related to 'Trump' were very common regardless of media PCA score. 'Barcelona' was also one of the most common words in tweets by both the Left and the Right. (There was a 17 August 2017 van attack in that city that killed 13 individuals (at the time of data collection) and injured more than 100 others.⁷) However, as we can see from the words in Table 4, there are also many differences in the

⁶We use stop words from the NLTK Python library, and our word counts also exclude the following words: 't', 'https', 'co', 'RT', 's', 'amp', 'n', 'w', and 'c'.

⁷One of those wounded individuals died later (after the time of data collection) from their injuries.

Table 4 The 25 most common words for nodes with negative (i.e., Left) and positive (i.e., Right) media PCA scores

Left		Right	
Word	Count	Word	Count
Charlottesville	98782	Charlottesville	84282
Trump	19352	Trump	11376
realDonaldTrump	10289	Obama	8195
white	9472	white	8174
Nazis	7743	DineshDSouza	8026
Nazi	6451	POTUS	7614
comments	5068	pastormarkburns	7394
charlottesville	4759	Barcelona	6348
good	4693	supremacist	6004
people	4637	organizer	5864
response	4091	rally	5851
must	4080	violence	5671
hate	3933	guy	5501
Barcelona	3930	MSM	5070
violence	3884	hate	4882
supremacy	3642	Right	4531
introducing	3584	city	4490
attack	3460	Americans	4489
via	3382	larryelder	4421
RepCohen	3358	Antifa	4409
rally	3345	Since	4183
Impeachment	3341	11	4011
Articles	3326	Chicago	3946
Klansmen	3310	Statues	3905
right	3155	40	3900

The blue text indicates words that appear in the top-25 words for the Left but not for the Right, and the red text indicates words that appear in the top-25 words for the Right but not for the Left.

words that were used by the Left and the Right. We illustrate such differences by coloring the relevant words. For example, ‘Obama’ was the third-most common word in tweets that were sent by nodes with positive media PCA scores, but it was not in the top one hundred for nodes with negative media PCA scores. Additionally, ‘Nazi’ appeared commonly in tweets from the Left, but it did not appear often in tweets from the Right. By contrast, the words ‘Antifa’ and ‘MSM’ were used often by the Right but not by the Left.

We also observe qualitative differences between the tweet content of the Left and Right in their shared common words, such as ‘Trump’ and ‘Barcelona’, in the #Charlottesville data set. The ‘Trump’ subset⁸ for which we have media PCA scores consists of 34,084 total tweets (of which about 32% are unique) from the Left and 18,791 total tweets (of which about 23% are unique) from the Right.⁹ As we show in the left set of columns of Table 5, the Left and Right conversations about Trump differ markedly from each other. The ‘Barcelona’ subset consists of 4779 tweets (of which 1672 are unique) from the Left and 7669 tweets (of which 1401 are unique) from the Right. In the right set of columns of Table 5, we show the 20 most common words for the Barcelona subset for both Left

⁸Aside from ‘Charlottesville’, which comes directly from the hashtag that we used to generate the data set, the word ‘Trump’ was the most common word for both Left and Right.

⁹Duplicate tweets arise, for example, from retweets. It is also possible that multiple accounts independently posted identical content (in addition to retweets).

Table 5 The 20 most common words (excluding words that correspond to handles of non-verified accounts) from the subset of tweets that include (left set of columns) the word ‘Trump’ and (right set of columns) the word ‘Barcelona’

Left		Right		Left		Right	
Word	Count	Word	Count	Word	Count	Word	Count
Charlottesville	30995	Charlottesville	16218	Charlottesville	4213	Charlottesville	7384
Trump	19396	Trump	11376	Barcelona	3926	Barcelona	6345
realDonaldTrump	10291	realDonaldTrump	3245	attack	1023	Muslims	2101
Nazis	5600	MAGA	1403	Trump	1020	13	2036
comments	4941	POTUS	1278	terrorism	842	right	2035
good	3788	President	1229	2	665	CNN	2019
introducing	3580	Romney	1209	prayers	450	kills	2006
Impeachment	3331	comments	1175	realDonaldTrump	438	condemned	2002
white	3318	apologize	1164	thoughts	420	kill	1996
Articles	3313	racist	1149	Terror	415	someone	1977
RepCohen	3306	blame	1096	gets	400	attack	1707
Klansmen	3301	Mayor	1085	settle	399	left	1686
must	2699	antifa	980	directed	399	lunatic	1632
Congress	2606	charlottesville	959	scolding	399	One	1621
censure	2540	Vice	958	intentional	397	wholesale	1620
supremacy	2239	Barcelona	880	ambigu	395	johncardillo	1617
wake	2107	left	868	condemns	356	copycat	1190
defense	2060	alt	818	took	346	Trump	718
NancyPelosi	2054	coming	781	comment	317	Blitzer	648
repulsive	2018	non	776	immediately	317	people	642

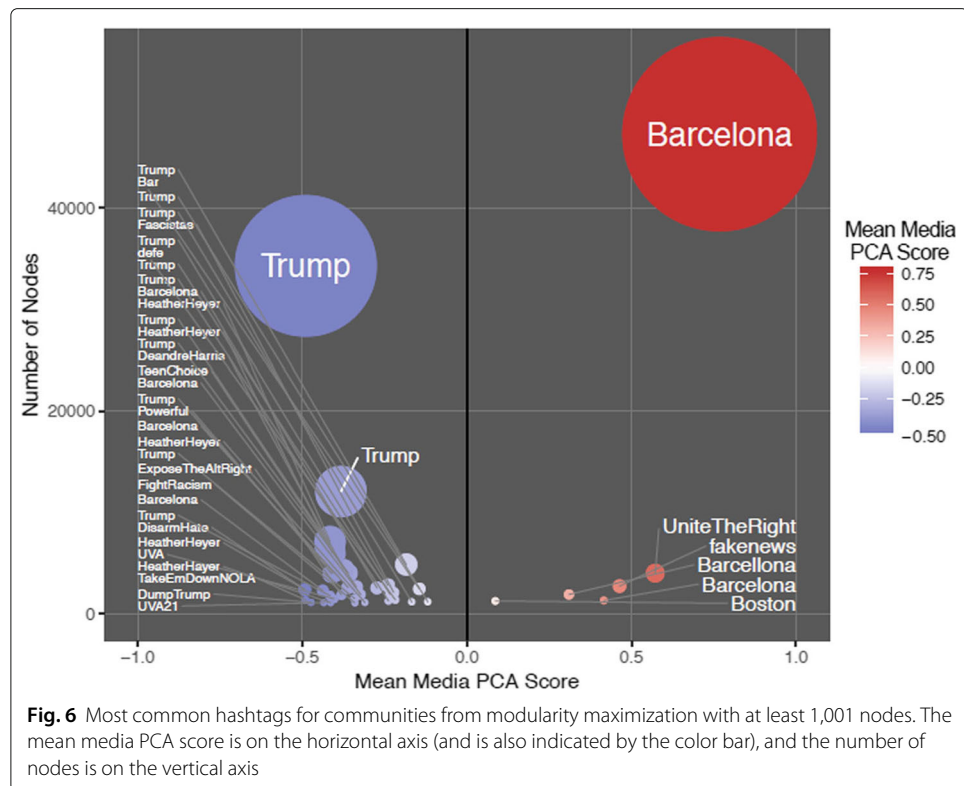
The blue text indicates words that appear in the top-20 words for the Left but not for the Right, and the red text indicates words that appear in the top-20 words for the Right but not for the Left.

and Right. For the Right, our examination of the most heavily retweeted tweets suggests that much of the discussion about ‘Barcelona’ in our data set involves comparing media coverage of the Charlottesville and Barcelona attacks. On the Left, some of the heavily retweeted tweets about ‘Barcelona’ centered on comparing Trump’s reaction to the Barcelona and Charlottesville attacks.

We also implement an idea from Gentzkow and Shapiro (2010), who used a chi-square statistic to analyze the different phrase usage of Democrats and Republicans in Congressional speeches. We apply their approach to words in tweets from the Left and Right in the ‘Trump’ subset (specifically, using equation (1) in Gentzkow and Shapiro (2010) with ‘phrases’ that consist of a single word) and find that the five words (which include ‘Nazis’, ‘antifa’, and ‘Vice’) with the largest chi-square values are also among the most common words in the ‘Trump’ subset (see Table 5). Therefore, we observe some consistency in results across different methods.

We also use hashtags to compare tweets between the Left and Right communities. In Fig. 6, we show the most common hashtag for each community,¹⁰ together with the community’s mean media PCA score. On the Left, the most common hashtags are #Trump (in 13 of 35 communities), #HeatherHeyer (in 5 of 35 communities, if we include a single community with ‘#HeatherHayer’), and #Barcelona (in 4 of 35 communities). Other top hashtags include #ExposeTheAltRight, #DumpTrump, #FightRacism, and #DisarmHate. On the Right, #Barcelona is the most common hashtag in 3 of 6 communities (if we

¹⁰This neglects hashtags that include ‘Charlottesville’, as our data collection was based on the #Charlottesville hashtag.



include a single community with #Barcellona). Other top hashtags on the Right are #UniteTheRight (from a community with an account of large in-degree whose Twitter handle references an influential account that identifies with the alt-right (O'Brien 2018)) and #fakenews (from a community with an account of large in-degree whose handle references a well-known white-supremacist hate symbol).

The fractions of tweets with unique content also differ between the Left and Right. Nodes with negative media PCA scores posted a total of 112,314 tweets, of which 42,458 (i.e., about 38%) are unique. Nodes with positive media PCA scores posted a total of 92,575 tweets, of which 22,462 (i.e., about 24%) are unique. We also observe a larger fraction of original content in the Left than in the Right when restricting to several specific topics, including 'Trump' (32% for the Left versus 23% for the Right), 'Barcelona' (35% for the Left versus 18% for the Right), 'MSM' (21% for the Left versus 8% for the Right), 'Obama' (31% for the Left versus 7% for the Right), and 'Antifa' (50% for the Left versus 22% for the Right). However, the proportion of tweets with the word 'Nazi' that are unique is slightly larger on the Right (21% for the Left versus 26% for the Right).

Conclusions and discussion

Our study of the Twitter conversation about #Charlottesville illustrates that (1) one can reasonably characterize nodes in terms of their media followership and a one-dimensional PCA-based Left/Right orientation score, (2) the Charlottesville retweet network is highly polarized with respect to this measure of Left/Right orientation, and (3) communities in the retweet network are largely homogeneous in their Left/Right node composition. Our

findings thus indicate that, with a few exceptions, the Twitter conversation largely split along ideological lines, as measured by the media preference of Twitter accounts.

As we just summarized, our investigation illustrates strong polarization in the Twitter conversation about #Charlottesville. We found that media followership on Twitter is informative and that the #Charlottesville retweet network is strongly assortative with respect to a PCA-based Left/Right orientation score. Our finding of positive assortativity with respect to media preference on Twitter is also consistent with previous studies of Twitter data (Colleoni et al. 2014; Conover et al. 2011; Feller et al. 2011; Pennacchiotti and Popescu 2011). However, in contrast with these previous studies, our approach to node characterization does not require text analysis or labeled data for training. Characterizing nodes via a principal component analysis of media followership is simple, easy to interpret, and provides a valuable complement to characterizing nodes based on the content of their tweets. Because the #Charlottesville retweet network is strongly assortative with respect to media preference, it is a potentially useful indicator of marked polarization on Twitter about the 'Unite the Right' rally and its aftermath. However, whether differences in media preferences are a cause or an effect (or both) of assortativity on social media is not something that our approach allows us to conclude.

Polarization is also evident in the community structure of the retweet network, as the communities are highly segregated in terms of their Left/Right node composition. The Left has a larger proportion of tweets with original content (as opposed to retweets) than the Right, and nodes with large hub scores tended to retweet nodes on the Left instead of those on the Right. We also found that modularity maximization detects Left communities with central nodes from disparate focal areas (such as business, media, entertainment, and politics). A robust feature of our community-detection results is that there are more large Left-leaning communities than Right-leaning ones. We also found that Left nodes with large in-degree have larger ratios of inter-community in-degree to intra-community in-degree, suggesting that heavily retweeted nodes on the Left were more likely to be retweeted by communities other than their own. Taken together, these findings suggest that Twitter accounts on the Left that condemned the 'Unite the Right' rally in Charlottesville and Trump's handling of the aftermath came from broad segments of society. Support on the Right in our data was concentrated in fewer communities, the largest of which includes the five most heavily retweeted accounts on the Right: FoxNews and right-wing personas DineshDSouza, pastormarkburns, larryelder, and johncardillo.

An important limitation of our study is that Twitter users are not a representative sample of the general population (Mitchell et al. 2014), and hashtag sampling introduces its own set of biases (Tüfekci 2014). Moreover, differences in Twitter usage and propensity to tweet political content may also differ with political affiliation (Colleoni et al. 2014). Consequently, it is important to compare our findings from Twitter to offline information. Our findings are consistent with a Quinnipiac poll that suggested that nearly one third of Republicans (but only 4% of Democrats) considered counterprotesters to be more to blame than white supremacists for the violence at Charlottesville (Quinnipiac University 2017). Several of the communities on the Right that we obtained from modularity maximization of the #Charlottesville retweet network also appear to reflect core participants of the 'Unite the Right' rally (Fausset and Feuer 2017), as indicated by the referencing

by central nodes in these communities of white-supremacist hate symbols or influential personalities in the alt-right.

Our investigation illustrated a stark distinction between Left and Right when we examined tweets that include the word 'Trump' with criticism on the Left versus support on the Right (see Table 5). For example, the most common hashtags from the Left in tweets that include 'Trump' were #Impeachment and #ImpeachTrump; by contrast, the most common hashtags were #Barcelona, #MAGA, and #fakenews from accounts in communities on the Right. Our findings are consistent with the extreme polarization and political tribalism in United States society that have been described by other studies (Dimrock and Carroll 2014; Iyengar and Westwood 2015; Jacobson 2017). Such societal divisions are apparent on Twitter, as documented both by the present study and by prior ones (Colleoni et al. 2014; Conover et al. 2011; Feller et al. 2011; Pennacchiotti and Popescu 2011), including recent work that suggested that polarization on Twitter is increasing over time (Garimella and Weber 2017).

It is also important to examine the role that fully automated accounts ('bots') and partially automated accounts (which have been dubbed 'cyborgs' (Chu et al. 2012)) play in shaping conversations (especially political ones) on Twitter and other social media platforms (Beskow and Carley 2018; Bessi and Ferrara 2016; Chu et al. 2012; Davis et al. 2016; Freitas et al. 2015; Stella et al. 2018; Yardi et al. 2010). Although an in-depth analysis of the role of bots in the #Charlottesville discussion is beyond the scope of the present paper, it is likely that many bot accounts are present in our data set. For example, it has been noted that automated naming schemes are an indicator of bot accounts (Beskow and Carley 2018). Naming schemes that end in sequences of eight digits, as well as accounts that consist of hexadecimal strings, both occur in the #Charlottesville data. Detailed investigation of these accounts and their behavior is an important topic for future work. Sockpuppet accounts (i.e., false accounts that are operated by an entity (Yamak et al. 2018)), such as those that are operated by the Internet Research Agency in St. Petersburg, Russia (House Permanent Select Committee on Intelligence Minority Staff 2017; United States District Court for the District of Columbia 2018), can also play important roles in content propagation and thus warrant further investigation. Antipathy and distrust across party lines can provide opportunities for actors who seek to fan societal divisions. For example, our data set includes tweets by prominent accounts that were operated by the Internet Research Agency (House Permanent Select Committee on Intelligence Minority Staff 2017; United States District Court for the District of Columbia 2018) and attacked both the Left and the Right.

Our approach of using media choice and PCA to characterize nodes is simple; it does not require labeled data, and it can also be used to study other fascinating topics. For example, it would be interesting to apply our approach to examine polarization on other topics (e.g., Brexit) and to see how polarization across political divides and attempts to bridge them change over time. It is not clear whether engagement with Twitter accounts with different viewpoints will decrease or increase polarization on divisive topics. For example, the empirical results of Bail et al. (2018) suggest that exposure on Twitter to contrasting ideologies can lead to increased polarization. An interesting question is how exposure shapes viewpoints of individuals with 'centrist' media preferences or ideologies. Our investigation focused primarily on the sign of a media PCA score, but the underlying media PCA score is continuous, and one can use it to examine media preferences

in a more nuanced way. In particular, characterizing accounts with moderate ('centrist') media PCA scores, studying network structure and tweet content by these accounts, and tracking the evolution of these characteristics over time are both feasible and relevant.

It is possible to refine our approach in various ways. For example, Landgraf and Lee (2015) used a type of PCA that is built for Boolean data, and it may be insightful to compare our PCA results to ones that employ their approach. Examining additional PCA components besides the first one (on which we focused exclusively in the present paper) may also be helpful for understanding the diversity of interactions between media and other accounts on Twitter. Careful analysis of additional principal components, interpretations of those components, and examining their relationship to network structure merit further investigation.

Examining multiple ideological dimensions (e.g., as in studies of voting by legislators on bills (Lewis et al. 2019)) and simultaneous analysis of multiple types of Twitter relationships using the formalism of multilayer networks (Kivelä et al. 2014) may also help deepen understanding of political communities and their news preferences after landmark news events. More broadly, we expect that our approach is generalizable to other contexts, and it may be helpful for examining other types of node characterizations (such as by analyzing different media outlets or types of followed accounts). Comparing network relationships and propagation of content across multiple social media platforms is of particular interest, as amount, diversity, and characteristics of platform usage vary across different segments of the world's population. There have been studies of the propagation of memes (Zannettou et al. 2018), web addresses (Zannettou et al. 2019), and anti-semitic content (Finkelstein et al. 2018) across different social media platforms; and it is important to undertake further studies of linkages across networks and their effects on recruitment, content propagation, and public discourse.

Abbreviations

API: Application program interface; HITS: Hyperlink-induced topic search; PCA: Principal components analysis; UTC: Coordinated universal time

Acknowledgements

We thank Rob Bond, Heather Brooks, Matt Salganik, and Sam Zhang for helpful comments.

Authors' contributions

JHT, MCE, STC, and MAP designed the study. JHT performed the analyses. JHT, MCE, STC, and MAP wrote the paper. All authors read and approved the final manuscript.

Funding

The authors have no funding sources to acknowledge for this study.

Availability of data and materials

We have made the de-identified data for the largest weakly connected component of the retweet network, as well as code for analyzing these data, available (at <https://osf.io/487fw/>) via the Open Science Framework. For privacy reasons, we do not provide any account names; for similar reasons, we do not provide tweet content.

Competing interests

The authors declare that they have no competing interests.

Author details

¹Department of Mathematics and Mathematical Biosciences Institute; The Ohio State University, Columbus OH, USA.

²Department of Epidemiology, Center for the Study of Complex Systems, and Department of Mathematics; University of Michigan, Ann Arbor MI, USA.

³Precision Health Enterprise and Institute for Next Generation Healthcare, Mount Sinai Health System, New York NY, USA.

⁴Department of Mathematics; UCLA, Los Angeles CA, USA.

Received: 21 May 2019 Accepted: 18 October 2019

Published online: 31 January 2020

References

- Allcott H, Gentzkow M (2017) Social media and fake news in the 2016 election. *J Econ Perspect* 31(2):211–236
- Arenas A, Duch J, Fernández A, Gómez S (2007) Size reduction of complex networks preserving modularity. *New J Phys* 9(6):176
- Baek EC, Porter MA, Parkinson C (2019) Social network analysis for social neuroscientists. arXiv:1909.11894
- Bail CA, Argyle LP, Brown TW, Bumpus JP, Chen H, Fallin Hunzaker M.B., Lee J, Mann M, Merhout F, Volfovsky A (2018) Exposure to opposing views on social media can increase political polarization. *Proc Nat Acad Sci USA* 115(37):9216–9221
- Barberá P, Wang N, Bonneau R, Jost JT, Nagler J, Tucker J, González-Bailón S (2015) The critical periphery in the growth of social protests. *PLoS ONE* 10(11):e0143611
- Beguerisse-Díaz M, Garduño-Hernández G, Vangelov B, Yaliraki SN, Barahona M (2014) Interest communities and flow roles in directed networks: The Twitter network of the UK riots. *J Royal Soc Interface* 11(101):20140940
- Benigni MC, Joseph K, Carley KM (2017) Online extremism and the communities that sustain it: Detecting the ISIS supporting community on Twitter. *PLoS ONE* 12(12):e0181405
- Beskow DM, Carley KM (2018) It's all in a name: Detecting and labeling bots by their name. arXiv:1812.05932
- Bessi A, Ferrara E (2016) Social bots distort the 2016 U. S. Presidential election online discussion. *First Monday* 21(11). <http://dx.doi.org/10.5210/fm.v21i11.7090>
- Blondel VD, Guillaume J-L, Lambiotte R, Lefebvre E (2008) Fast unfolding of communities in large networks. *J Stat Mech Theory Exp* 2008(10):10008
- Bougoussa M, Ben Romdhane L (2015) Identifying authorities in online communities. *ACM Trans Intell Syst Technol* 6(3):30
- Brin S, Page L (1998) The anatomy of a large-scale hypertextual Web search engine. *Comput Netw ISDN Syst* 30(1–7):107–117
- Bryden J, Funk S, Jansen VAA (2013) Word usage mirrors community structure in the online social network Twitter. *EPJ Data Sci* 2(1):3
- Budak C, Goel S, Rao JM (2016) Fair and balanced? Quantifying media bias through crowdsourced content analysis. *Publ Opin Quarter* 80(51):250–271
- Catania JA, Coates TJ, Kegels S, Fullilove MT, Peterson J, Marin B, Siegel D, Hulley S (1992) Condom use in multiethnic neighborhoods of San Francisco — The population-based AMEN (AIDS in Multiethnic Neighborhoods) study. *Am J Publ Health* 82(2):284–287
- Cha M, Haddadi H, Benevenuto F, Gummadi KP (2010) Measuring user influence in Twitter: The million follower fallacy. In: 4th International AAAI Conference on Weblogs and Social Media (ICWSM 2010). pp 10–17
- Chew C, Eysenbach G (2010) Pandemics in the age of Twitter: Content analysis of Tweets during the 2009 H1N1 outbreak. *PLoS ONE* 5(11):e14118
- Chu Z, Gianvecchio S, Wang H, Jajodia S (2012) Detecting automation of Twitter accounts: Are you a human, bot, or cyborg? *IEEE Trans Depend Sec Comput* 9(6):811–824
- Cihon P, Yasserli T (2016) A biased review of biases in Twitter studies on political collective action. *Front Phys* 4:34
- Colleoni E, Rozza A, Arvidsson A (2014) Echo chamber or public sphere? Predicting political orientation and measuring political homophily in Twitter using big data. *J Commun* 64(2):317–332
- Conover MD, Ratkiewicz J, Francisco M, Gonçalves B, Flammini A, Menczer F (2011) Political polarization on Twitter. In: Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media (ICWSM 2011). pp 89–96
- Curbelo C (2017) We've already fought this war. *South Dade News Leader*. (17 August 2017). http://www.southdadenewsleader.com/opinion/we-ve-already-fought-this-war/article_ac44ce12-83b7-11e7-9a4c-0f8cc879c501.html
- Daniels J (2018) The algorithmic rise of the “alt-right”. *Contexts* 17(1):60–65
- Davis CA, Varol O, Ferrara E, Flammini A, Menczer F (2016) BotOrNot: A system to evaluate social bots. In: Proceedings of the 25th International Conference on World Wide Web — Companion Volume (WWW 2016 Companion). pp 273–274
- Denecke K, Kriek M, Otrusina L, Smrz P, Dolog P, Nejdil W, Velasco E (2013) How to exploit Twitter for public health monitoring? *Methods Inf Med* 52(4):326–339
- Dimrock M, Carroll D (2014) Political polarization in the American public: How increasing ideological uniformity and partisan antipathy affect politics, compromise, and everyday life. Tech Rep. <http://www.people-press.org/2014/06/12/political-polarization-in-the-american-public/>
- Duggan P, Jouvenal J (2019) Neo-Nazi sympathizer pleads guilty to federal hate crimes for plowing car into protesters at Charlottesville rally. https://www.washingtonpost.com/local/public-safety/neo-nazi-sympathizer-pleads-guilty-to-federal-hate-crimes-for-plowing-car-into-crowd-of-protesters-at-unite-the-right-rally-in-charlottesville/2019/03/27/2b947c32-50ab-11e9-8d28-f5149e5a2fda_story.html. Accessed 11 Nov 2019
- Edler D, Rosvall M (2019) The MapEquation software package. <http://www.mapequation.org>. Accessed 13 Feb 2019
- Fausset R, Feuer A (2017) Far-right groups surge into national view in Charlottesville. *New York Times*. <https://www.nytimes.com/2017/08/13/us/far-right-groups-blaze-into-national-view-in-charlottesville.html>. Accessed 11 Nov 2019
- Feller A, Kuhnert M, Sprenger TO, Welpel I (2011) Divided they tweet: The network structure of political microbloggers and discussion topics. In: Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media (ICWSM 2011). pp 474–477
- Fenn DJ, Porter MA, Mucha PJ, McDonald M, Williams S, Johnson NF, Jones NS (2012) Dynamical clustering of exchange rates. *Quant Finan* 12(10):1493–1520
- Finkelstein J, Zannettou S, Bradlyn B, Blackburn J (2018) A quantitative approach to understanding online antisemitism. arXiv:1809.10644
- Flaxman S, Goel S, Rao JM (2016) Filter bubbles, echo chambers, and online news consumption. *Publ Opin Quarter* 80(51):298–320
- Fortunato S, Barthélemy M (2007) Resolution limit in community detection. *Proc Nat Acad Sci USA* 104(1):36–41
- Fortunato S, Hric D (2016) Community detection in networks: A user guide. *Phys Rep* 659:1–44
- Fosdick BK, Larremore DB, Nishimura J, Ugander J (2018) Configuring random graph models with fixed degree sequences. *SIAM Rev* 60(2):315–355

- Freelon D, McIlwain CD, Clar M (2016) Beyond the hashtags: #Ferguson, #Blacklivesmatter, and the online struggle for offline justice. <https://ssrn.com/abstract=2747066>
- Freeman L (1977) A set of measures of centrality based on betweenness. *Sociometry* 40(1):35–41
- Freitas C, Benevenuto F, Ghosh S, Veloso A (2015) Reverse engineering socialbot infiltration strategies in Twitter. In: Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM '15). pp 25–32
- Garimella K, Weber I (2017) A long-term analysis of polarization on Twitter. In: International AAAI Conference on Web and Social Media (ICWSM 2017). pp 528–531
- Gentzkow M, Shapiro JM (2010) What drives media slant? Evidence from U.S. daily newspapers. *Econometrica* 78(1):35–71
- Gentzkow M, Shapiro JM (2011) Ideological segregation online and offline. *Quarter J Econ* 126(4):1799–1839
- Ghalmame Z, Hassouni ME, Cherifi C, Cherifi H (2019) Centrality in modular networks. *EPJ Data Sci* 8:15
- Gleich DF (2015) PageRank beyond the Web. *SIAM Rev* 57(3):321–363
- Guimerà R, Amaral LAN (2005) Functional cartography of complex metabolic networks. *Nature* 433(7028):895–900
- Good BH, de Montjoye Y, Clauset A (2010) Performance of modularity maximization in practical contexts. *Phys Rev E* 81(4):046106
- Groseclose T, Milroy J (2005) A measure of media bias. *Quarter J Econ* 120(4):1191–1237
- House Permanent Select Committee on Intelligence Minority Staff (2017) HPSCI Minority Exhibit B. https://democrats-intelligence.house.gov/uploadedfiles/exhibit_b.pdf. https://democrats-intelligence.house.gov/uploadedfiles/exhibit_b.pdf
- Huber LP (2009) Challenging racist nativist framing: Acknowledging the community cultural wealth of undocumented Chicana college students to reframe the immigration debate. *Harvard Education Rev* 79(4):704–730
- Iyengar S, Westwood SJ (2015) Fear and loathing across party lines: New evidence on group polarization. *Am J Polit Sci* 59(3):690–707
- Jacobson GC (2017) The triumph of polarized partisanship in 2016: Donald Trump's improbable victory. *Polit Sci Quarter* 132(1):9–41
- Jeub LGS, Balachandran P, Porter MA, Mucha PJ, Mahoney MW (2015) Think locally, act locally: Detection of small, medium-sized, and large communities in large networks. *Phys Rev E* 91(1):012821
- Jeub LGS, Bazzi M, Jutla IS, Mucha PJ (2011–2016) A generalized Louvain method for community detection implemented in Matlab. Version 2.0. <https://github.com/GenLouvain/GenLouvain>
- Jolliffe IT (2002) *Principal Components Analysis*. 2nd. Springer-Verlag, Heidelberg, Germany
- Kivela M, Arenas A, Barthelemy M, Gleeson JP, Moreno Y, Porter MA (2014) Multilayer networks. *J Compl Netw* 2(3):203–271
- Kleinberg JM (1999) Authoritative sources in a hyperlinked environment. *J Assoc Comput Mach* 46(5):604–632
- Lakoff G, Ferguson S (2006) The framing of immigration. <https://escholarship.org/uc/item/0j89f85g>. Accessed 11 Nov 2019
- Lambiotte R, Delvenne J-C, Barahona M (2015) Random walks, Markov processes and the multiscale modular organization of complex networks. *IEEE Trans Netw Sci Eng* 1(2):76–90
- Landgraf AJ, Lee Y (2015) Dimensionality reduction for binary data through the projection of natural parameters. [arXiv:1510.06112](https://arxiv.org/abs/1510.06112)
- Lee K, Mahmud J, Chen J, Zhou M, Nichols J (2014) Who will retweet this?: Automatically identifying and engaging strangers on Twitter to spread information. In: Proceedings of the 19th International Conference on Intelligent User Interfaces (IUI '14). pp 247–256
- Leicht EA, Newman MEJ (2008) Community structure in directed networks. *Phys Rev Lett* 100(11):118703
- Lerman K, Ghosh R (2010) Information contagion: An empirical study of the spread of news on Digg and Twitter social networks. In: Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media (ICWSM 2010). pp 90–97
- Lewis JB, Poole KT, Rosenthal H (2019) Voteview. <https://voteview.com>. Accessed 11 Nov 2019
- Lynch M, Freelon D, Aday S (2014) Syria in the Arab Spring: The integration of Syria's conflict with the Arab uprisings, 2011–2013. *Res Politics* 1(3):2053168014549091
- McNeill A, Harris PR, Briggs P (2016) Twitter influence on UK vaccination and antiviral uptake during the 2009 H1N1 pandemic. *Front Publ Health* 4:26
- Mitchell A, Gottfried J, Kiley J, Masa KE (2014) Political polarization & media habits. Tech Rep. <http://www.journalism.org/2014/10/21/political-polarization-media-habits/>
- Mitchell A, Hitlin P (2014) Twitter reaction to events often at odds with overall public opinion. Tech Rep. <http://www.pewresearch.org/2013/03/04/twitter-reaction-to-events-often-at-odds-with-overall-public-opinion/>
- Morales AJ, Borondo J, Losada JC, Benito RM (2015) Measuring political polarization: Twitter shows the two sides of Venezuela. *Chaos* 25(3):033114
- Morgan J (2017) How the 'alt-right' came to dominate the comments on Trump's Facebook page. *The Atlantic*. <https://www.theatlantic.com/politics/archive/2017/01/how-the-alt-right-influenced-trump-supporters-language-on-facebook/513593/>. Accessed 11 Nov 2019
- Morstatter F, Pfeffer J, Liu H, Carley KM (2014) Is the sample good enough? Comparing data from Twitter's streaming API with Twitter's firehose. In: Proceedings of the Seventh International AAAI Conference on Weblogs and Social Media (ICWSM 2014). pp 400–408
- Mucha PJ, Richardson T, Macon K, Porter MA, Onnela J-P (2010) Community structure in time-dependent, multiscale, and multiplex networks. *Science* 328(5980):876–878
- Newman MEJ (2002) Assortative mixing in networks. *Phys Rev Lett* 89(20):208701
- Newman MEJ (2003) Mixing patterns in networks. *Phys Rev E* 67(2):026126
- Newman MEJ (2006) Finding community structure in networks using the eigenvectors of matrices. *Phys Rev E* 74(3):036104
- Newman MEJ (2018) *Networks*. 2nd edition. Oxford University Press, Oxford, UK
- Newman MEJ, Girvan M (2004) Finding and evaluating community structure in networks. *Phys Rev E* 69(2):026113

- O'Brien L (2018) Trump's most influential white nationalist troll is a Middlebury grad who lives in Manhattan. https://www.huffpost.com/entry/trump-white-nationalist-troll-ricky-vaughn_n_5ac53167e4b09ef3b2432627
- Otero V (2018) Media Bias Chart: Version 4.0. <https://www.adfontesmedia.com>
- Pariser E (2011) *The Filter Bubble: What the Internet Is Hiding from You*. Penguin Random House UK, London, UK
- Park J, Barabási A-L (2007) Distribution of node characteristics in complex networks. *Proc Natl Acad Sci USA* 104(46):17916–17920
- Peixoto TP (2017) Bayesian stochastic blockmodeling. In: Doreian P, Batagelj V, Ferligoj A (eds). *Advances in Network Clustering and Blockmodeling*. John Wiley & Sons, Inc., New York City, NY, USA. arXiv:1705.10225 (in press)
- Pennacchiotti M, Popescu A (2011) A machine learning approach to Twitter user classification. In: *Proceedings of AAAI Conference on Weblogs and Social Media (ICWSM 2011)*. pp 281–288
- Perrin A (2015) Social media usage: 2005–2015. <http://www.pewinternet.org/2015/10/08/social-networking-usage-2005-2015/>. Accessed 11 Nov 2019
- Porter MA, Onnela J-P, Mucha PJ (2009) Communities in networks. *Notices Am Math Soc* 56(9):1082–1097, 1164–1166
- Quinnipiac University (2017) Trump is dividing the country, U.S. voters say 2-1, Quinnipiac University national poll finds; most trust media more than President. (23 August 2017). <https://poll.qu.edu/national/release-detail?ReleaseID=2482>
- Romero DM, Galuba W, Asur S, Huberman BA (2011) Influence and passivity in social media. In: *Proceedings of the 20th International Conference on World Wide Web — Companion Volume (WWW 2011 Companion)*. pp 113–114
- Rosvall M, Bergstrom CT (2008) Maps of random walks on complex networks reveal community structure. *Proc Natl Acad Sci USA* 105(4):1118–1123
- Salathé M, Jones JH (2010) Dynamics and control of diseases in networks with community structure. *PLoS Comput Biol* 6(4):e1000736
- Salathé M, Vu DW, Khandelwal S, Hunter DR (2013) The dynamics of health behavior sentiments on a large online social network. *EPJ Data Sci* 2(1):4
- Salganik MJ (2017) *Bit by Bit: Social Research in the Digital Age*. Princeton University Press, Princeton, NJ, USA
- Shannon CE (1948) A mathematical theory of communication. *Bell Syst Tech J* 27(3):379–423
- Shear MD, Haberman M (2017) Trump defends initial remarks on Charlottesville: Again blames 'both sides'. *New York Times*. <https://www.nytimes.com/2017/08/15/us/politics/trump-press-conference-charlottesville.html>. Accessed 11 Nov 2019
- Signorini A, Segre AM, Polgreen PM (2011) The use of Twitter to track levels of disease activity and public concern in the US during the influenza A H1N1 pandemic. *PLoS ONE* 6(5):e19467
- Stella M, Ferrara E, De Domenico M (2018) Bots increase exposure to negative and inflammatory content in online social systems. *Proc Natl Acad Sci USA* 115(49):12435–12440
- Sunstein CR (2001) *Republic.com*. Princeton University Press, Princeton, NJ, USA
- Towers S, Afzal S, Bernal G, Bliss N, Brown S, Espinoza B, Jackson J, Judson-Garcia J, Khan M, Lin M, Mamada R, Moreno VM, Nazari F, Okuneye K, Ross ML, Rodriguez C, Medlock J, Ebert D, Castillo-Chavez C (2015) Mass media and the contagion of fear: The case of Ebola in America. *PLoS ONE* 10(6):e0129179
- Tremayne M (2014) Anatomy of protest in the digital era: A network analysis of Twitter and Occupy Wall Street. *Soc Mov Stud* 13(1):110–126
- Tüfekçi Z (2014) Big questions for social media big data: Representativeness, validity and other methodological pitfalls. In: *Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media (ICWSM 2014)*. pp 505–514
- Tüfekçi Z (2017) *Twitter and Tear Gas: The Power and Fragility of Networked Protest*. Yale University Press, New Haven, CT, USA
- Tüfekçi Z, Wilson C (2012) Social media and the decision to participate in political protest: Observations from Tahrir Square. *J Commun* 62(2):363–379
- Twitter, Inc. (2019) About verified accounts. <https://help.twitter.com/en/managing-your-account/about-twitter-verified-accounts>. Accessed 8 May 2019
- United States District Court for the District of Columbia (2018) *United States of America v. Internet Research Agency LLC*, Case 1:18-cr-00032-DLF. 18 U.S.C. 2, 371, 1349, 1028A. <https://www.justice.gov/file/1035477/>. Accessed 11 Nov 2019
- Weng J, Lim E-P, Jiang J, He Q (2010) TwitterRank: Finding topic-sensitive influential Twitterers. In: *Proceedings of the 3rd ACM International Conference on Web Search and Data Mining (WSDM 2010)*. pp 261–270
- Weng L, Menczer F, Ahn Y-Y (2013) Virality prediction and community structure in social networks. *Sci Rep* 3:2522
- Wojcieszak M, Mutz D (2009) Online groups and political discourse: Do online discussion spaces facilitate exposure to political disagreement? *J Commun* 59(1):40–56
- Wong FMF, Tan CW, Sen S, Chiang M (2016) Quantifying political leaning from tweets, retweets, and retweeters. *IEEE Trans Knowl Data Eng* 28(8):2158–2172
- Yamak Z, Saunier J, Vercouter L (2018) SocksCatch: Automatic detection and grouping of sockpuppets in social media. *Knowledge-Based Systems* 149:124–142
- Yardi S, Romero D, Schoenebeck G, Boyd D (2010) Detecting spam in a Twitter network. *First Monday* 15(1). <https://doi.org/10.5210/fm.v15i1.2793>
- Zannettou S, Caulfield T, Blackburn J, De Cristofaro E, Sirivianos M, Stringhini G, Suarez-Tangil G (2018) On the origins of memes by means of fringe Web communities. In: *Proceedings of the ACM Internet Measurement Conference 2018 (ACM IMC 2018)*. pp 188–202
- Zannettou S, Caulfield T, Setzer W, Sirivianos M, Stringhini G, Blackburn J (2019) Who let the trolls out?: Towards understanding state-sponsored trolls. In: *Proceedings of the 10th ACM Conference on Web Science (WebSci '19)*. pp 353–362

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.