Characterizing the Twitter network of prominent politicians and SPLC-defined hate groups in the 2016 US presidential election

Raazesh Sainudiin¹ · Kumar Yogeeswaran² · Kyle Nash³ · Rania Sahioun²

Received: 21 December 2018 / Revised: 22 May 2019 / Accepted: 25 May 2019 / Published online: 19 July 2019
© The Author(s) 2019

Abstract

We characterize the Twitter networks of the major presidential candidates, Donald J. Trump and Hillary R. Clinton, with various American hate groups defined by the US Southern Poverty Law Center (SPLC). We further examined the Twitter networks for Bernie Sanders, Ted Cruz, and Paul Ryan, for 9 weeks around the 2016 election (4 weeks prior to the election and 4 weeks post-election). We carefully account for the observed heterogeneity in the Twitter activity levels across individuals through the null hypothesis of apathetic retweeting that is formalized as a random network model based on the directed, multi-edged, self-looped, configuration model. Our data revealed via a generalized Fisher's exact test that there were significantly many Twitter accounts linked to SPLC-defined hate groups belonging to seven ideologies (Anti-Government, Anti-Immigrant, Anti-LGBT, Anti-Muslim, Alt-Right, White-Nationalist and Neo-Nazi) and also to @realDonaldTrump relative to the accounts of the other four politicians. The exact hypothesis test uses Apache Spark's distributed sort and join algorithms to produce independent samples in a fully scalable way from the null model. Additionally, by exploring the empirical Twitter network we found that significantly more individuals had the fewest retweet degrees of separation simultaneously from Trump and each one of these seven hateful ideologies relative to the other four politicians. We conduct this exploration via a geometric model of the observed retweet network, distributed vertex programs in Spark's GraphX library and a visual summary through neighbor-joined population retweet ideological trees. Remarkably, less than 5% of individuals had three or fewer retweet degrees of separation simultaneously from Trump and one of several hateful ideologies relative to the other four politicians. Taken together, these findings suggest that Trump may have indeed possessed unique appeal to individuals drawn to hateful ideologies; however, such individuals constituted a small fraction of the sampled population.

Keywords Donald Trump · Twitter · 2016 US presidential election · US hate groups · Configuration model · Scalable generalized Fisher's exact test · Apache Spark · Directed degrees of separation · Empirical geometric retweet model · Population retweet ideological trees

✉ Raazesh Sainudiin
raazesh.sainudiin@math.uu.se

Kumar Yogeeswaran
kumar.yogeeswaran@canterbury.ac.nz

Kyle Nash
knash@ualberta.ca

Rania Sahioun
rouna79@yahoo.com

¹ Department of Mathematics, Combient Competence Centre for Data Engineering Sciences, Uppsala University, Uppsala, Sweden

² Department of Psychology, University of Canterbury, Christchurch, New Zealand

³ Department of Psychology, University of Alberta, Edmonton, Canada

1 Introduction

The 2016 US presidential election will be remembered as one of the most divisive in recent history with two of the least liked candidates competing for the White House (Enten 2016; ABC News/Washington Post 2016). During the election, the mainstream media appeared baffled by the rise of Donald J. Trump, a businessman with no prior experience in government. Two narratives emerged to explain Trump's rise to fame. On the one hand, it was argued that Trump's rise was driven by uncertainty and angst in the American public (Pew Research Center 2015; Reicher and Haslam 2017; Hochschild 2016) as well as frustration and distrust of a partisan political system (Packer 2016; Roussos 2016; Ball 2016). This first narrative broadly reflects a number

of psychological theories that commonly demonstrate that uncertainty and angst can catalyze the formation of rigid, nationalistic, and extreme political beliefs; these are supported by research on right-wing authoritarianism (Altemeyer 1996), uncertainty–identity theory (Hogg 2007), need for closure (Webster and Kruglanski 1994), and motivated social cognition (Jost et al. 2003). By contrast, a second narrative argued that Donald Trump’s rise was driven by a growing sense of threat among White Americans regarding their changing status in America as a result of growing cultural diversity and the improving status of minority groups made salient by the election of Barack Obama. Such changes were argued to promote minority scapegoating and a campaign that pitted “us” versus “them” (Webster and Kruglanski 1994; Jost et al. 2003; Kharakh and Primack 2016; Knowles and Tropp 2016; Beauchamp 2017; Ingraham 2016; Anderson 2016). Along with uncertainty–identity theory (Hogg 2007), this second narrative is reflected in psychological research on relative deprivation theory (Runciman 1966; Walker and Pettigrew 1984), intergroup threat theory (Stephan and Stephan 2000), and recent evidence that ethnic diversity can threaten the status, identity, and distinctiveness of majority group members (Major et al. 2017; Craig and Richeson 2014; Yogeeswaran and Dasgupta 2014).

Though both sets of factors may have collectively played a role in Trump’s rise to power, there has been considerable debate within the media and around dinner tables about the latter narrative that Trump was able to capitalize on White Americans’ fears about their changing status and social identity in an unstable national landscape. Critics of Donald Trump point toward Trump’s own rhetoric on the campaign trail, which often featured insults leveled against Mexicans, Muslims, immigrants, and women (Hananoki 2016; Shalby 2016). Such rhetoric, it was argued, took advantage of the increased appetite for nationalism and outgroup derogation due to perceived threats to status and identity for the majority group. By doing so, Trump was argued to have forged a coalition where normative bigotry was thought to be acceptable (Crandall and White 2016). As a consequence, Trump’s candidacy was openly supported by the Ku Klux Klan (KKK) and Neo-Nazi groups (Reuters 2016; Raymond and Pyke 2017). Though Trump’s campaign was careful to distance itself from the endorsement of such extreme hate groups (Detrow 2016), critics point to the rise of hate crimes since the election (Hayden 2017; Piggott 2017; Southern Poverty Law Center 2017; Müller and Schwarz 2018) and the Trump administration’s decision to remove government focus on right-wing extremism as evidence of Trump’s ties to hate groups (O’Connor 2017; Raymond and Pyke 2017). Even months after the election, people continued to debate the role of bigotry in Trump’s candidacy (Wood 2017; Tuttle 2017; Martin 2017).

Twitter has become a major platform for communication between politicians and their followers. Consequently, Twitter activity has been used to gauge political sentiments, predict elections (Tumasjan et al. 2010; Vaccari et al. 2015) and reportedly influence the 2016 election (Cadwalladr 2017; Grassegger and Krogerus 2017; Woolley and Guibeault 2017).¹ President Trump uses the platform to directly communicate with the American public (Selk 2017; Editors 2017). Twitter has also become a major avenue through which hate groups spread their ideologies with larger sections of the population (Hale 2012). For all these reasons, Twitter was chosen as the ideal medium to observe naturalistic relationships between various individuals and groups.

In the present work, we examined the Twitter networks of several key politicians and hate groups or their leadership during the 2016 US presidential election to shed light on two major narratives regarding Trump’s rise to fame. As most academic studies in the Twitterverse, we were limited by the freely available Twitter data. So our study of a specifically sampled population, despite our attempts at optimizing the sampling design, may not be representative of the US Twitter population. This caveat should be borne in mind when generalizing this study without the full data that can be purchased from Twitter.

The political leaders included in our study were the two major presidential candidates (i.e., Hillary Clinton with mentionable screen-name as @HillaryClinton and Donald Trump as @realDonaldTrump). Additionally, we chose to include Ted Cruz (as @tedcruz) and Bernie Sanders (as @BernieSanders) who were the second most popular Republican and Democratic presidential candidates in the 2016 election. And finally, we chose to include Paul Ryan (as @SpeakerRyan) as he was the Speaker of the House of Representatives at the time of the election. These individuals all possess active Twitter accounts with more than 2 million followers and engage in frequent Twitter activity. We collected publicly streaming data from these five politicians’ official Twitter accounts.

To examine the links between these five US politicians and hateful ideologies, we used the Southern Poverty Law Center’s (SPLC’s) definition and ideological classification of hate groups and hate group leaders; hereafter referred to as hate groups. Although the SPLC has been criticized for including groups and individuals with politically diverging views from their own (Freilich 2006; Torres 2017; Jonsson 2013; Graham 2016), it is the most comprehensive database for hate groups in the USA that is available to the public with records over the last 40 years (Gilliard-Matthews 2011;

¹ This article is the subject of separate legal complaints on behalf of Cambridge Analytica LLC and SCL Elections Limited, and Sophie Schmidt.

Chermak et al. 2013; Ryan and Leeson 2011; Gemignani and Hernandez-Albujar 2015). Moreover, our mathematical models, statistical tests, and the data science pipelines are generic and can therefore be used to examine Twitter networks between political leaders and any database of hate groups as long as these have enough activity on Twitter. More generally, our approach formally only requires a population of users, and two sub-populations of interest.

The SPLC does not consider all groups or individuals on its list as violent or engaged in criminal activities, but rather identifies any group or individual “whose beliefs or practices attack or malign an entire class of people, typically for their immutable characteristics.” (Southern Poverty Law Center 2016). The database does not include foreign hate groups or extremist groups such as ISIS, Al Qaeda, or Boko Haram, as the focus of the database is on American hate groups. In the present work, we focused only on hate groups which the SPLC has specifically designated as promoting a particular hateful ideology in order to examine the appeal of major politicians for subscribers to that ideology. We collected publicly streaming data from 52 Twitter accounts belonging to eight hateful ideologies including: Neo-Nazi or NN (e.g., National Socialist Movement), Black-Separatist or BIS (e.g., Nation of Islam), Anti-LGBT or AL (e.g., Westboro Baptist Church), Anti-Immigrant or AI (e.g., American Border Patrol), Anti-Muslim or AM (e.g., ACT for America), Alt-Right or AR (e.g., American Renaissance), Anti-Government or AG (e.g., Oath Keepers), and White-Nationalist or WN (e.g., Aryan Brotherhood). As argued by others (Glaser et al. 2002), studying blatant bigotry and hate is better suited using online communication as opposed to more traditional means of conducting prejudice research via self-report measures due to social norms against such blatant expressions. In line with this, the present work examines linkages between hate groups and politicians by focusing on social media communication on Twitter that offers such hate groups a shroud of anonymity.

We use the *retweet network*, which represents each retweet as a directed edge from the tweeter, i.e., the user who posted the original tweet to the retweeter, i.e., the user who made the retweet, to model the communication links between users in our dataset. A retweet not only represents the strongest index of interest in the message, but also endorsement and trust in the communicator (Jansen et al. 2009; Metaxas et al. 2015). Retweet analysis via random network models (as used in this study) also circumvents the ambiguity and further uncertainty associated with statistical algorithms in natural language understanding (NLU) that will be required when working with (a) quoted tweets, where one is allowed to add a comment to the retweet, (b) reply tweets, or (c) reply of quoted tweets, where one can reply in possible disagreement, etc. Thus, retweets are a simple and natural signal of directional concurrence, i.e., concurrence

with the user who posted the (original) tweet. This is especially the case when one retweets another multiple times. Examining retweet networks would therefore be ideal as the research questions of interest focus on observing linkages between people endorsing blatantly hateful ideologies and endorsing specific political candidates during the 2016 Presidential election.

Using the observed retweet network, we examine three specific questions: (Q1) Is Trump preferentially retweeted by hate groups when compared to other politicians against the null random network model of *apathetic retweeting*? (Q2) What frequency of unique users retweeted both a politician and a hate group more than one would expect under the null model? And (Q3) What is the joint distribution of the *degrees of separation*, measured through the lengths of the most retweeted directed paths in the observed Twitter network, to each user from each of the five politicians and the eight most prolific hateful ideologies on Twitter? The last question is of a data-exploratory nature, as such a joint distribution can give insights into certain projections of ideological profiles of politically active Twitter users involving *population ideological trees* that naturally extend typical projections involving a binary spectrum, where users are distinguished only up to points along an interval from one extreme to another.

Our analyses observed Twitter networks without consciously filtering out bot accounts. As bots may have played a role in influencing the 2016 election (Woolley and Guibeault 2017; Howard et al. 2016), we simply observed the public Twitter streams associated with these politicians and hate groups. However, we do retrospectively study the effect of potential bot accounts with high entropy scores and the accounts identified by Twitter as being Russian trolls (Twitter Public Policy 2018) as discussed later on.

Using exact tests, we address questions (Q1) and (Q2) by attempting to reject the null hypotheses of *apathetic retweeting* where each user is apathetic (i.e., indifferent with a lack of interest or concern²) about who they retweet while preserving the number of times they retweet others and the number of times others retweet them. We formalize this null model as the *apathetic retweet network* from the directed versions of the configuration network model (Newman et al. 2001; Bollobás 2001) [see Aldous (2013) and references therein for an overview of models for stochastic social dynamics]. Thus, we can preserve the observed innate differences among users’ tweet and retweet rates in this null model when we attempt to reject it in favor of its complementary alternative hypothesis of non-apathetic retweeting of one set of users by another set of users. By leveraging

² “Apathy.” Merriam-Webster.com. Accessed July 10, 2018. <https://www.merriam-webster.com/dictionary/apathy>.

Apache Spark (Zaharia et al. 2016), a unified engine for big data processing, our scalable and fault-tolerant version of CUTPERMUTEANDREWIRE, the original Monte Carlo sampling algorithm due to Newman et al. (2001), can produce independent samples from this null model. These samples can in turn be used to obtain the null distribution with consistent estimates of the p value of any test statistic of the apathetic retweet network, for instance. Thus, we have an exact non-parametric test that can computationally scale for arbitrarily large directed networks. Finally, to answer (Q3) we use distributed vertex programs over the empirically estimated *geometric retweet network* model to obtain the degrees of separation from a politician and a hate group and visualize the pairwise distance in the degrees of separation between users through a neighbor-joined retweet ideological tree of the sampled population.

There is a large body of literature on online social media and hate. In general, such studies are either descriptive [see, for e.g., Bliuc et al. (2018) and the references therein and Mondal et al. (2017) for a recent example] or predictive [for e.g., hate speech detection (Schmidt and Wiegand 2017; Fortuna and Nunes 2018)]. Thus, the quantitative literature around social media and hate is more focused on mining and modeling patterns in the data for predictive and descriptive purposes, as opposed to testing specific hypotheses of interest to social scientists using custom-designed experiments. This work is a significant extension of Sainudin et al. (2018) that primarily focused on testing specific hypotheses. Here, we not only focus on classical statistical hypothesis testing with interpretable probabilistic models of the null hypothesis, but also on a systematic exploration of the data through scalable data science pipelines for conducting (designing, testing, and exploring) one's own statistical experiments in Twitter by leveraging the latest advances in elementary distributed algorithms for sorting, joining, and vertex programming.

2 Data, design, models, and methods

2.1 Data and experimental design

2.1.1 Communications in Twitter

Twitter, a popular micro-blogging service, provides an observable social network with millions of users (Jansen et al. 2009). Twitter allows users to communicate or update their status in many ways. One can post a message called a *tweet* that is no more than 140 characters in length (in 2016), *follow* another Twitter user, and receive the status updates of every user they follow. A tweet that is shared publicly with one's own followers is known as a *retweet*.

2.1.2 Dataset

Our dataset comprises of 21,749,868 communication events in Twitter over a total of 9 weeks centered around the US 2016 election. The tweet IDs are available from <https://tinyurl.com/y8lamxxz>. The dataset was collected in our designed experiment using Twitter's streaming and REST APIs (Developers 2018) that were extended for Apache Spark (Zaharia et al. 2016). Each of these communication events was parsed and classified into seven types of events using SparkSQL modules.

Streaming data collector Over 17 million events were collected from Twitter's public streams by directly tracking communications related to the Twitter accounts of the five political candidates and 52 hate group accounts. Because only 78% of the SPLC identified hate groups had a valid Twitter account that we could track, our study is limited to a further subset of 52 user accounts who were active in the Twitter public streams. Thus, our approach is not exhaustive, in terms of being able to track every account of each hateful ideology, but is nonetheless representative of the public Twitter activity.³

Retrospective data augmentation The remaining 5 million events were obtained retrospectively using Twitter's REST API as follows. Due to most retweets being an immediate reaction to a tweet that one finds interesting or concurs with Kwak et al. (2010), nearly 7 million of the 10.5 million retweets in the 17 million events collected from the public streams happened within the same day of the original tweet and over 98% of the retweets happened within a week of the original tweet. Furthermore, since our tweet collector is only recording events that are directly related to one of our tracked users, we do not know of any other Twitter interactions by those who retweeted one of our tracked users. Thus, to better understand the recent retweet behavior of at least some of the politically active Twitter users with other non-tracked users on Twitter, we focus on October 19 2016, the day of the 3rd US Presidential debate, and obtain a seed set of users who retweeted either @HillaryClinton or @realDonaldTrump on this day. The communication intensity reached over 120 events per second around the debate in our streaming data collector. Our seed set is made up of a random sample of about a third of all users (including

³ Although the SPLC also identifies ideologies of the Ku Klux Klan (KKK), Neo-confederates, Racist Skinheads, and Holocaust Deniers, within their database, there was not enough activity on the Twitter accounts of these groups that we were able to track given just a few thousand followers. It may be that these groups have a smaller online presence, or they simply use a different social media platform [including closed discussion forums; see Martin (2017)] making it harder to follow as we focused our analyses only on hate groups with public streams.

all verified and geo-enabled accounts) who retweeted either Clinton or Trump on October 19 2016 and thus constitutes an evenly represented sample of politically active Twitter users from the two parties. For each user in the seed set, at the end of the 9-week period, we added all the retweets from their 200 most recent status updates that occurred in the 9-week period. This strategy involves a breadth-first expansion about the seed set of users in the much larger retweet network on Twitter as it allows us to expand our 9-week-long retweet network by focusing on the recent retweet timelines of those who retweeted either of the two final presidential candidates during the day of the last debate. Crucially, these augmented data added another 0.3 million users to our network, increased the number of retweet events from 10.5 million to 13.7 million, and made the retweet network into a single connected component.

2.2 Modeling retweets

2.2.1 Retweet network

Retweets are a simple and natural signal of *directional concurrence*, i.e., concurrence *with* the user who posted the (original) tweet *by* the user who retweeted it (i.e., between the tweeter and the retweeter), as they express interest in the message and also endorsement and trust in the communicator (Jansen et al. 2009; Metaxas et al. 2015), especially when retweeted multiple times. Retweet analysis via random network models as done in this study circumvents the ambiguity and uncertainty associated with statistical algorithms (Barberá et al. 2015) in natural language understanding (NLU) that will be required when working with (a) *quoted tweets* where one is allowed to add a comment to the retweet, (b) *reply tweets* or (c) *reply of quoted tweets* where one can reply in possible disagreement.

The number of retweets per day captured from the public stream during the 9-week period reached over 1 million during the day of the third US Presidential debate and the days leading up to the election. The proportion of retweets is known to vary greatly depending on the features of the Twitter subnetwork under study (Kwak et al. 2010). There were over 13.7 million retweets (63.1%) of an original tweet and over 2.7 million (12.6%) original tweets in our dataset. We ignore the remaining 25% of the events that require further NLU and focus our analysis instead on the retweet network obtained from 75% of all events in our dataset as explained in the next section.

We allowed the set of users who tweet and retweet, i.e., tweeters and retweeters, to form the nodes of the retweet network. Each retweet was allowed to represent a directed edge or arc from its tweeter to its retweeter, i.e., the tweeter–retweeter pair, as a signal of directional concurrence in the retweet network. Because one can retweet more than

Table 1 Retweet network statistics of the five political accounts

Politician	in-degree	in-nbhd	out-degree	out-nbhd
Donald Trump	40	12	5,952,257	958,262
Hillary Clinton	225	121	2,774,111	943,995
Bernie Sanders	107	62	762,209	356,718
Paul Ryan	769	158	68,973	28,902
Ted Cruz	322	189	49,479	27,663

one tweet posted by any user including oneself, we allowed for parallel or multiple edges between the same pair of users including self-loops, i.e., edges from and to the same user. There were more than 4.4 million unique tweeter–retweeter pairs out of 2.5 million unique users representing over 16.4 million tweets and retweets in our dataset with over 21.7 million communication events during the 9-week period around the US presidential election. Mathematically, the retweets in our dataset are represented by a *directed multi-edged self-looped network*. Thus, the *out-degree* (number of outgoing edges) and *in-degree* (number of incoming edges) for a user in the retweet network gives the number of times that the user is retweeted by others and the number of times the user retweets others (including oneself), respectively. Similarly, the number of distinct users who are retweeted by the user is given by their *in-nbhd* and the number of distinct retweeters of the user is given by their *out-nbhd*.

The observed retweet network is highly heterogeneous and largely dominated by the two presidential candidates as depicted in Table 1. Due to our Twitter collector’s design with retrospective data augmentation, our network forms a single connected component when viewed with undirected edges.

2.2.2 Apathetic retweet network model

The *configuration model* for directed networks (Newman et al. 2001) is a random network model that produces samples uniformly from the set of networks that preserve the in-degree and out-degree of each node in given observed network.

By representing each retweet as a directed edge between the author of the tweet and its retweeter, we define the null apathetic model for a random retweet network as the *directed, multi-edged, self-looped configuration model*, whereby each user apathetically or non-preferentially retweets another user, by choosing uniformly at random from the set of all users, while preserving their observed in-degree and out-degree. However, the model frees up information on the users involved in retweeting, i.e., who retweets whom, provided the in-degree and out-degree of each user is identical to that of the observed network. This null model thus preserves the innate tweet and retweet rates of each user

while modeling complete apathy or indifference for who one chooses to retweet.

2.2.3 Cut, permute and rewire

CUTPERMUTEANDREWIRE is a Monte Carlo algorithm to generate independently sampled networks from the directed multi-edged self-looped random configuration model, i.e., our null random network model for apathetic retweeting. CUTPERMUTEANDREWIRE is a distributed, scalable, and fault-tolerant version of the standard construction involving random pairings of out-bound and in-bound half-edges (Newman et al. 2001) through the following three steps: (i) *cutting* the directed edges representing the retweets in our observed retweet network into out-bound and in-bound half-edges, (ii) *permuting* the in-bound half-edges by sorting them according to pseudo-random numbers that are generated and associated with them, and (iii) *rewiring* the original out-bound half-edges with the permuted in-bound half-edges using a distributed join. Note that the in-degree and out-degree of each node in the observed retweet network is preserved after the three steps by construction.

By taking advantage of the fastest available distributed sorting and optimized distributed joining algorithms (Zaharia et al. 2016), CUTPERMUTEANDREWIRE can produce independent sample networks with tens of millions of retweets or edges in a small Apache Spark cluster over six commodity compute nodes. The null distribution of any test statistic of the retweet network can be directly obtained from applying it to each independent Monte Carlo sample from our scalable fault-tolerant randomized algorithm with probability given by:

$$1/(\text{number of retweets})!$$

under the null model. By comparing the observed test statistic to Monte Carlo samples from the null distribution, one can directly obtain consistent estimates of the p value in order to attempt to reject the null hypothesis of apathetic retweeting in favor of the alternative hypothesis of non-apathetic retweeting in the framework of generalized Fisher's exact test.

2.2.4 Degrees of separation in geometric retweet network model

To gain deeper insights into the empirical retweet network beyond the immediate retweet neighborhood of each user, i.e., those retweeted by a user and those who retweet the user, and explore the alternative hypothesis space if the null hypothesis is rejected, we adapt Milgram's concept of *six degrees of separation* (Milgram 1967) that all people in the world are six or fewer steps away from each other so that a sequence of "a friend of a friend" relationships can be made

to connect any two people in a maximum of six steps. We adapt the concept in three major ways.

First, we focus on Twitter users and replace the mutual or undirected relationship of being a friend by the directed relationship of being a retweeter. This adaptation accounts for the main difference between various social and technological networks (Watts and Strogatz 1998) as well as other communication networks (Leskovec and Horvitz 2008) that are characterized by mutually reciprocal relationships, and the directed relationships in Twitter where a path from a user to another may follow several distinct sequences while not existing in the reverse direction (Kwak et al. 2010).

Second, instead of the degrees of separation between every pair of users in an undirected sense, we are interested in the *retweet* degrees of separation in a directional sense, measured by the length of the most retweeted path that originates from a few influential users, such as politicians or hate groups, and terminates at each user.

Third, we account for the strength of the retweet relationship when defining the most retweeted path by incorporating $r_{a,b}$, the observed number of retweets between the user who is the source of the original tweets, i.e., the tweeter a , and the user who retweets them, i.e., the retweeter b , through the directed edge-weight given by $\hat{p}_{a,b} = 1/(1 + r_{a,b})$ that is used to specify the probability of an independent geometric random variable giving the number of retweets of user a by user b .

Let G be such a weighted directed retweet network with nodes as users and directed edges with weights:

$$\{\hat{p}_{a,b} = (1 + r_{a,b})^{-1} : \forall \text{ edge } (a, b) \text{ in } G \text{ with } r_{a,b} > 0\}.$$

The collection of independent but non-identical geometric random variables with probability parameters given by these weights is our empirically estimated *geometric retweet network* model for the joint distribution of the number of retweets of each user by another in our dataset.

We derive our estimated geometric retweet network model by recalling the well-known relationship between Poisson, exponential, and geometric random variables. If the random variable $R_{a,b}$ giving the total number of retweets of a by b is Poisson distributed with a random mean parameter $\xi_{a,b}$ that is drawn from the exponentially distributed random variable with rate parameter $1/\lambda_{a,b}$, then $R_{a,b}$ is geometrically distributed with probability parameter $p_{a,b} = 1/(1 + \lambda_{a,b})$ and expectation $\lambda_{a,b}$. We can estimate the parameters from the observed number of retweets $r_{a,b}$ via the moment estimate $\hat{\lambda}_{a,b} = r_{a,b}$ and model the number of retweets during the 9-week period according to the geometric random variable with probability parameter $\hat{p}_{a,b} = (1 + r_{a,b})^{-1}$ for each directed edge between a tweeter and a retweeter in the retweet network.

A small weight $\hat{p}_{a,b} = (1 + r_{a,b})^{-1}$ corresponds in an inversely proportionate manner to a large number of

Table 2 Relative frequency of retweets by any one of the hate groups for any original tweet made by one of the politicians

Politician				
Donald Trump	Hillary Clinton	Bernie Sanders	Paul Ryan	Ted Cruz
0.987 : (0.6008, 0.6013)	0 : (0.2708, 0.2709)	0 : (0.0677, 0.0682)	0 : (0.00411, 0.00413)	0.0131 : (0.0024, 0.0028)

Observed test statistic: marginal interval for the region of acceptance at 0.001 significance level

retweets, and the shortest path from user a to user b through this directed weighted network corresponds to the path with a large number of retweets from user a to user b . Thus, the estimated geometric model interpretation of the weighted retweet network G allows us to use a straightforward distributed vertex program via Pregel API in Apache Spark's GraphX library to obtain the shortest path that is composed of a sequence of d tweeter–retweeter pairs of edges, say,

$$E = \{(a, v_1), (v_1, v_2), \dots, (v_{d-1}, b)\},$$

with the lowest sum of weights given by $\sum_{(u,v) \in E} \hat{p}_{u,v}$ among all possible paths between an influential user of interest a and every other user b in G . Crucially, this $\hat{p}$ -weighted shortest path is called *the most retweeted path* as it is composed of the same sequence of edges with the correspondingly high sum of retweet counts given by $\sum_{(u,v) \in E} r_{u,v}$. The length d of the most retweeted path from a to b , known as the (*retweet*) *degrees of separation* from a to b , has a clear interpretation as the length of the sequence of “retweeter of a retweeter” statements along the most retweeted path that is needed to link user a to user b by considering the retweet activities of every user in the network. When we have a set of influential users $A = \{a_1, a_2, \dots, a_k\}$ of interest, say k accounts of a hateful ideology, we define the shortest path from A to any user b as the minimum of the k shortest paths from each user in A to b .

3 Results and discussion

Given the observed heterogeneity in the activity levels of Twitter users, including the most influential user @realDonaldTrump, who was more active than the other four politicians with over twice as many retweets by about the same number of retweeters as the next most influential user @HillaryClinton (Table 1), any hypothesis test needs to control for this heterogeneity. Our approach uses retweet networks sampled independently from the null apathetic retweet network model in order to obtain the null distribution of any test statistic while preserving the observed differences in the activity levels, in terms of the observed in-degree and out-degree of each user in the retweet network. Note that we do not have multiple testing issues when attempting to reject the null hypothesis of apathetic retweeting in favor of its complementary alternative hypothesis of non-apatetic

retweeting since we obtain independent Monte Carlo samples from the joint distribution of the statistics under the null model to estimate the acceptance and rejection regions at the significance level of 0.001.

In the following two sections, we conducted two hypothesis tests involving frequencies of direct retweets and that of retweeters of politicians and hate groups to address our questions (Q1) and (Q2). Then, we see if the tests are robust to any effects by bot and/or troll accounts and in the last section we explore the empirical retweet network.

3.1 Frequencies of direct retweets of politicians by hate groups

None of the 194,098 retweets of an original tweet by one of the hate group accounts were made by any one of our five political accounts. However, out of the 7233 retweets made by one of the hate group accounts, 151 were retweets of Trump, 2 were retweets of Cruz, and none were of the remaining three politicians. These 151 retweets of Trump were split among NN (87), WN (55), AM (6), and AG (3) groups, while the two retweets of Cruz were from NN and WN hate groups.

One of the simplest statistics is the relative frequency of retweets by any one of the hate groups of any original tweet made by each of the five politicians. This statistic of relative frequencies is given by five proportions that sum to 1. Given that Trump has more than twice as many retweets as Clinton despite having about the same number of retweeters (see out-degree and out-nbhd in Table 1), one may expect Trump to be retweeted more by hate groups even if they were retweeting the five politicians apathetically without showing any preference for Trump. To control for this effect, we used the CUTPERMUTEANDREWIRE algorithm to obtain samples from the joint distribution of the relative frequencies under the null hypothesis of apathetic retweeting. As shown in Table 2, the observed statistic lies outside the acceptance region obtained from 1000 Monte Carlo samples from the null model and we thus reject the null hypothesis of apathetic retweeting with a p value less than 0.001 in favor of non-apatetic retweeting with over 98% of retweets by hate groups preferring Trump well above the upper bound of 60.13% under the null model, while simultaneously avoiding any retweeting of Clinton or Sanders with a relative frequency of 0% that is well below their respective

Table 3 Observed frequency of distinct users who retweeted a politician and an account within a hate group at least five times each

Ideology	Politician				
	Donald Trump	Hillary Clinton	Bernie Sanders	Paul Ryan	Ted Cruz
Anti-Government (AG)	*107 : (0, 1)	3 : (0, 3)	0 : (0, 1)	*2 : (0, 1)	*4 : (0, 1)
Anti-Immigrant (AI)	*2314 : (375, 498)	°44 : (373, 492)	°15 : (369, 485)	*204 : (47, 95)	*133 : (18, 54)
Anti-LGBT (AL)	*121 : (0, 4)	1 : (0, 4)	1 : (0, 4)	*5 : (0, 3)	*23 : (0, 3)
Anti-Muslim (AM)	*215 : (0, 3)	0 : (0, 3)	0 : (0, 3)	*13 : (0, 3)	*21 : (0, 3)
Neo-Nazi (NN)	*45 : (0, 1)	0 : (0, 1)	0 : (0, 1)	0 : (0, 1)	0 : (0, 1)
White-Nationalist (WN)	*548 : (0, 12)	9 : (0, 10)	0 : (0, 10)	6 : (0, 8)	4 : (0, 7)
Black-Separatist (BIS)	°69 : (653, 811)	°40 : (649, 808)	°22 : (645, 801)	°5 : (72, 128)	°1 : (28, 66)
Alt-Right (AR)	*22 : (0, 0)	0 : (0, 0)	0 : (0, 0)	0 : (0, 0)	*3 : (0, 0)

*, °Observed test statistic: marginal interval for the region of acceptance at 0.001 significance level

Fig. 1 Number of unique users who retweeted a politician and a hate group at least five times each. (Note: The y-axis is in log-scale in powers of 10)

lower bounds of 27.08% and 6.77% under the null model, respectively. The observed relative frequencies for Ryan and Cruz, on the other hand, deviate relatively little from the corresponding marginal intervals of their acceptance region.

3.2 Frequencies of retweeters of politicians and hate groups

For a robust test statistic, we looked one degree deeper into the retweet network and obtained the number of Twitter users who retweeted both a politician and a specific hate group at least five times. This provided a more conservative estimate of the number of users retweeting hate groups or politicians as it would only include those actively retweeting both. The observed frequency of such users is shown in Fig. 1 and Table 3 along with their 0.1% marginal intervals of the acceptance region for the null hypothesis of apathetic

retweeting. Since the observed frequencies jointly lie well outside the acceptance region obtained from 1000 Monte Carlo samples drawn from the null model using the CUT-PERMUTEANDREWIRE algorithm, we reject the null hypothesis of apathetic retweeting in favor of non-apathetic retweeting with a p value less than 0.001.

When the observed counts are in the rejection region by falling below or above the respective marginal intervals of the acceptance region, we can obtain statistically meaningful insights of the observations under the favored alternative hypothesis of non-apathetic retweeting. However, such insights need further models and tests for a refined and rigorous understanding within the space of alternative hypotheses.

In Table 3, the observed frequencies that lie in the rejection region above and below the marginal intervals for the acceptance region are indicated by * and °, respectively.

Thus, for example, a statistically significant frequency of users have retweeted Trump, Ryan or Cruz and Anti-Immigrant groups at least five times each in a non-apathetic and preferentially approving manner since these frequencies are above their null intervals, i.e., the marginal intervals of the acceptance region for the null model of apathetic retweeting at a significance level of 0.001. These same frequencies are below the null intervals for Clinton and Sanders in favor of the non-apathetic retweeting alternative hypothesis, but in a preferentially disapproving manner. Only Trump has a significant frequency of retweeters who also frequently retweet AG, AI, AL, AM, NN, WN, and AR hate groups in a non-apathetic and preferentially approving manner. Note that the Black-Separatist hate group forms a “natural control” as their retweeters non-apathetically retweet all five politicians in a preferentially disapproving manner.

Our results are robust to changes in the threshold number of retweets being at least 5, albeit at the expense of fewer observations for larger thresholds. This is because the pattern of user frequencies across hateful ideologies and the five politicians shown in Fig. 1 is preserved across other thresholds. The results are also robust when the test statistic is an expectation of the frequencies taken over all threshold levels greater than 1. We chose the naturally interpretable threshold of at least five retweets over the 9-week period of observation in Fig. 1 and Table 3 to select users with more than one retweet per fortnight, on average.

3.3 Effect of bot and troll accounts

While we did not actively filter out highly sophisticated bot accounts given the difficulty in their detection (Ferrara et al. 2016), we studied the effect of removing retweets related to accounts with very high scores for two types of information entropy based on: (i) the distributions of time lags between successive posts and (ii) the content of words within the posts. This is because accounts that send messages at uniform time intervals or post messages with unusually static or identical text content might be bots or cyborgs with unusual entropy scores. Accounts with more than 10 posts that crossed 2.5 standard deviations from the mean for either entropy score were flagged as potential bot accounts using a scalable variant of a practical approach (Kramer 2017). Additionally, we retrospectively examined our retweet network against the 2752 now-deactivated Twitter accounts that the company identified as being allegedly tied to Russia’s *Internet Research Agency* troll farm (Twitter Public Policy 2018). Out of the 12,984,331 retweets in our dataset, less than 0.1% were related to a troll account (293 were retweeted by and 12,347 were originally tweeted by a troll account), and out of 2,451,081 distinct users in our retweet network, only 172 were related to a troll account. Interestingly, removal of these troll-related retweets from the

retweet network did not alter the statistical tests in Tables 2 and 3. The observed test statistics remained the same when we removed the troll-related retweets and the acceptance region did not vary enough to alter any of the reported results. Similarly, removal of nearly 1% of all retweets that were affiliated with one of 2045 potential bot accounts or one of the troll accounts did not affect the test results. The robustness of the tests to presence or absence of the potential bot and identified troll accounts suggests that our results are independent of the direct influence of these bot and troll accounts, at least with regard to our Twitter study focused around the activities of the five politicians in relation to the 52 hate groups.

3.4 Degrees of separation from politicians and hate groups

For each user node in the estimated geometric retweet network model, we obtained the length of the most retweeted path, i.e., the degrees of separation, from one of the five politicians (DT = @realDonaldTrump, HC = @HillaryClinton, BS = @BernieSanders, PR = @SpeakerRyan, TC = @tedcruz) and eight hateful ideologies (AI = Anti-Immigrant, AM = Anti-Muslim, WN = White-Nationalist, AL = Anti-LGBT, AG = Anti-Govt, NN = Neo-Nazi, BIS = Black-Separatist, AR = Alt-Right). Thus, each user had a *retweet ideological profile* or simply a profile given by their thirteen degrees of separation (from the five politicians and eight hateful ideologies). In order to focus on users who were politically active, we excluded (i) unreachable users who had infinite degrees of separation from all five politicians, i.e., those users who were unreachable from any one of the five politicians by a sequence of retweets, and (ii) users whose most retweeted path did not have an average number of retweets greater than 4 per edge.

We thus obtained profiles for 2,137,712 politically active users over the 9-week period (our monitored population) and report in Table 4 the frequencies of users with the same retweet ideological profile given by thirteen numbers along with the corresponding percentage of the overall sample.

Using the Manhattan distance between profiles, i.e., sum of the absolute value of the differences between the profiles (i.e., the 13 degrees of separation), we clustered the top 15 groups of users with the most frequent profiles constituting 95% of the monitored population as shown by the classification tree (Saitou and Nei 1987; Boc et al. 2012) in the first column of Table 4. Since this tree is built from the most retweeted path-lengths from influential users promoting various ideologies, we call this the *retweet ideological tree* of the sampled population.

Broadly, there are three clades in the tree (only the top 15 most populous leaf nodes are shown). First, the politically ambivalent clade is made up of three profiles

Table 4 The top 15 groups of users according to their profiles of most retweeted path-lengths from the five politicians* and eight hateful ideologies° given by their id, frequency, percentage of population and their classification given by the retweet ideological tree with leaf nodes as the ids

Ideological tree id	Frequency	Percentage of population	Politician					Hate group							
			DT	HC	BS	PR	TC	AI	AM	WN	AL	AG	NN	BIS	AR
05	42,853	02.005	1	1	2	4	4	5	5	7	6	4	7	7	7
07	11,481	00.537	1	2	1	4	4	5	5	7	6	4	7	7	7
12	5868	00.274	1	1	1	4	4	5	5	7	6	4	7	7	7
11	5972	00.279	4	2	3	5	7	8	8	9	9	7	10	10	10
02	791,286	37.016	3	1	2	4	6	7	7	8	8	6	9	9	9
04	74,126	03.468	3	1	1	4	6	7	7	8	8	6	9	9	9
03	232,093	10.857	3	2	1	6	6	7	7	9	8	6	9	9	9
13	5173	00.242	3	1	1	6	6	7	7	8	8	6	9	9	9
01	811,586	37.965	1	4	7	4	4	5	5	7	6	4	7	7	7
15	3892	00.182	1	4	7	1	4	5	5	7	3	4	7	7	7
14	4011	00.188	1	4	7	4	4	1	5	3	5	4	5	7	7
08	10,460	00.489	3	5	9	1	3	3	3	5	3	6	7	9	9
09	8069	00.377	3	3	3	3	1	4	3	3	3	6	5	6	9
06	29,997	01.403	2	3	3	3	3	5	3	3	5	5	5	3	3
10	6257	00.293	1	3	3	4	4	5	3	3	5	4	5	3	3

*DT = @realDonaldTrump, HC = @HillaryClinton, BS = @BernieSanders, PR = @SpeakerRyan, TC = @tedcruz

°AI = Anti-Immigrant, AM = Anti-Muslim, WN = White-Nationalist, AL = Anti-LGBT, AG = Anti-Govt, NN = Neo-Nazi, BIS = Black-Separatist, AR = Alt-Right

(with leaf id's 05, 07, and 12) that constitute nearly 3% of the monitored population and is characterized by the smallest possible degrees of separation not only from @realDonaldTrump (DT), but also from either @HillaryClinton (HC) or @BernieSanders (BS). Second, the Democratic-leaning cluster that is characterized by having fewer degrees of separation from at least one of the two Democratic politicians: HC and BS. It is shown by the clade with leaf id's 02, 04, 03, 13, and the more diverged id 11 that is farthest from all hate groups and from the Politicians but relatively closer to HC. The three most populous leaf nodes in the Democratic-leaning clade with id's 02, 03, and 04, constitute about 37%, 11%, and 3.5% of the monitored population. Third, the Republican-leaning cluster that is characterized by having fewer degrees of separation from at least one of the three Republican politicians: @realDonaldTrump (DT), @SpeakerRyan (PR) and @tedcruz (TC). It is shown by the larger clade with the following seven leaf id's: 01, 15, 14, 08, 09, 06, and 10 in Table 4. The most populous leaf nodes in the Republican-leaning clade are made up of two of the 15 groups of users with id's 01 and 15, and constitute nearly 38% of the monitored population with distinct profiles that are primarily characterized by the smallest possible degrees of separation of 1 from DT. Similarly, the other two less populous Republican-leaning id's (08, 09, 06, and 10) are primarily characterized by different degrees

of separation from DT, PR, and TC, whereas id 14, with its own branch, is notably hateful by having the smallest possible degrees of separation from both DT and Anti-Immigrant groups.

Note that more than 91% and 95% of users have their degrees of separation from all eight hateful ideologies no smaller than 4 and 3, respectively. This is reflective of little direct influence by hate groups for over 95% of the monitored population. In order to get insights on users whose most retweeted paths from at least one of the hateful ideologies is within three or fewer retweet degrees of separation, we zoomed into the less frequent profiles of the remaining 5% of the monitored population of politically active users. These results are depicted as cumulative percentages of the monitored population for the following nine pairs of retweet degrees of separation from one of the five politicians and a given hateful ideology:

$\{(1, 1), (1, 2), (2, 1), (1, 3), (2, 2)(3, 1), (2, 3), (3, 2), (3, 3)\}$

in Fig. 2. For example, over 0.6% and 1% of the population is at 1 degree of separation from @realDonaldTrump and within 1 and 2 degrees of separation from Anti-Immigrant ideologies, as specified by the y-axis values corresponding to the x-axis values of (1, 1) and (1, 2), respectively, in the top left subplot of Fig. 2. Similarly, nearly 2.4% of the population is within 3 degrees of separation from Trump and Anti-Immigrant ideologies as specified by the y-axis

Fig. 2 Cumulative percentage of the monitored Twitter population who are within a given in-degree of separation from a politician and a hateful ideology

value corresponding to the x -axis value of (3, 3) in the same subplot. Thus, the cumulative percentage of the politically active monitored population that is within a given ordered

pair of degrees of separation (from a politician and a hateful ideology) is depicted similarly by the subplots of Fig. 2.

Although the number of distinct retweeters of @realDonaldTrump and @HillaryClinton is roughly the same at around 950,000 (see out-nbhd column in Table 1), the proportion of the population that is within 1 or 2 degrees of separation from @realDonaldTrump and from any one of the six hateful ideologies (AI, AM, AL, WN, AG, and NN), whose test statistic fell above the acceptance region and thus suggestive or indicative of non-apathetic preferential retweeting, is consistently higher than that from any other politician (unlike those for the Black-Separatists, our “natural control” in the exact test, and Alt-Right ideologies partly due to minimal Twitter activity induced by SPLC’s reclassification during the course of our data collection). As mentioned earlier, the null model of apathetic retweeting was strongly rejected using the frequency of users who retweeted both a politician and a hateful ideology more than four times each in favor of the alternative hypothesis of non-apathetic retweeting in which these same six ideologies were found to be retweeted in a preferentially approving manner by those who also retweeted Trump (Table 3). Thus, zooming into the less frequent profiles of the remaining 5% of the monitored population that is within three retweet degrees of separation from a politician and a hate group is a conservative exploration of the empirical geometric retweet network (see Fig. 2).

Recall that the degrees of separation has a clear interpretation as the length of the sequence of “retweeter of a retweeter” statements along the most retweeted path that links a set of influential users to each user by considering the global retweet activities of every user in the network. Thus, this exploratory analysis provides global insights into the most retweeted pathways from influential sets of users (politicians and leaders of hateful ideologies), through the joint distribution of the lengths of these paths, to each of the 2,137,712 politically active users. It further provides an ideological tree to classify users based on the Manhattan distance of these path-lengths that allows one to ask where one lies on a well-defined ideological tree, a mathematically natural generalization (Dress 1984; Dress and Terhalle 1996) of the line giving the usual bi-polar spectrum of left versus right or Democrat versus Republican, a common restriction in social media research (e.g., Barberá et al. 2015), and thus allowing us to consider the effects of transmissions of “memes” (Dawkins 1976, p. 192)⁴ on social media networks (Sainudiin and Welch 2016) through ideological trees using more sophisticated mathematical notions (Athreya et al. 2013, 2017).

4 Conclusion

Using 9 weeks of Twitter data collected around the 2016 US presidential election involving nearly 22 million communication events, the present research examined the Twitter linkages between five major American political leaders (Donald Trump, Hillary Clinton, Ted Cruz, Bernie Sanders, and Paul Ryan) with American hate groups. Using two different approaches to the data, we found converging evidence that Donald Trump possessed unique appeal to a variety of American hate groups.

First, utilizing direct retweets, we found through a generalized Fisher’s exact test that Twitter users who frequently retweeted hate groups (i.e., at least five times in the span of 9 weeks) were significantly more likely to retweet Trump over any other politician. A significant number of retweeters of nearly all hateful ideologies, except for Black-Separatists (BIS), were found in Trump’s Twitter network including Anti-Immigrant (AI), Anti-Muslim (AM), Anti-Government (AG), Neo-Nazi (NN), Alt-Right (AR), White-Nationalists (WN), and Anti-LGBT (AL). By comparison, retweeters of only AI, AM, and AL were significantly linked to Paul Ryan or Ted Cruz, but even these linkages were quite small relative to those observed in Trump’s network.

Second, using most retweeted path-lengths, we found that a significantly larger proportion of people were within one, two and three degrees of separation from Donald Trump and from AG, AL, AM, AI, WN, or NN ideologies, relative to the corresponding degrees of separation from Hillary Clinton, Bernie Sanders, Ted Cruz, or Paul Ryan and the same hate groups. None of the political candidates appeared to hold particular appeal to users retweeting BIS hate groups. While there is some debate about the relationship between hate groups and hate crimes (for e.g., Adamczyk et al. 2014; Green and Rich 1998; Ryan and Leeson 2011), it is still important to examine the social media networks of prominent politicians and hate groups as political leadership influences social norms about what is acceptable and what is not in wider society (Hogg and Reid 2006).

In the present work, our retweet network provides for the second narrative we mentioned in the introduction that Trump held unique appeal to those espousing hateful ideologies more so than other Republican or Democratic candidates. Nevertheless, it is important to note that only a small fraction of those retweeting Trump were within a few degrees of separation from hateful ideologies, suggesting that most of Trump’s support appear removed from people with hateful views.

Acknowledgements Open access funding provided by Uppsala University. R. Sainudiin was partly supported by the chaire Modélisation Mathématique et Biodiversité of Veolia Environnement–École Polytechnique–Museum National d’Histoire Naturelle–Fondation X.

⁴ See Solon (2013).

Distributed cloud computing of Twitter data by R. Sainudiin was supported by Databricks Academic Partners Program.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- ABC News/Washington Post (2016) Clinton hits a new high in unpopularity; on par with Trump among registered voters. Post poll, Langer Research Associates. <http://www.langerresearch.com/wp-content/uploads/1144-59ClintonTrumpFavorability.pdf>. Accessed 28 May 2017
- Adamczyk A, Gruenwald J, Chermak SM, Freilich JD (2014) The relationship between hate groups and far-right ideological violence. *J Contemp Crim Justice* 30(3):310–332
- Aldous D (2013) Interacting particle systems as stochastic social dynamics. *Bernoulli* 19(4):1122–1149
- Altemeyer B (1996) *The authoritarian specter*. Harvard University Press, Harvard
- Anderson C (2016) Donald Trump is the result of white rage, not economic anxiety. *Time*. November 16 2016. <http://time.com/4573307/donald-trump-white-rage/>. Accessed 28 May 2017
- Athreya S, Eckhoff M, Winter A (2013) Brownian motion on $\mathbb{R}$ -trees. *Trans Am Math Soc* 365(6):3115–3150
- Athreya S, Löhr W, Winter A et al (2017) Invariance principle for variable speed random walks on trees. *Ann Probab* 45(2):625–667
- Ball M (2016) Donald Trump and the politics of fear. *The Atlantic*. September 2 2016. <https://www.theatlantic.com/politics/archive/2016/09/donald-trump-and-the-politics-of-fear/498116/>. Accessed on 28 May 2017
- Barberá P, Jost JT, Nagler J, Tucker JA, Bonneau R (2015) Tweeting from left to right. *Psychol Sci* 26(10):1531–1542
- Beauchamp Z (2017) White riot. *Vox*. THE BLOG 10/21/2016 07:15 am ET | Updated October 21, 2016. <https://www.vox.com/2016/9/19/12933072/far-right-white-riot-trump-brexite>. Accessed 28 May 2017
- Bliuc AM, Faulkner N, Jakubowicz A, McGarty C (2018) Online networks of racial hate: a systematic review of 10 years of research on cyber-racism. *Comput Hum Behav* 87:75–86. <https://doi.org/10.1016/j.chb.2018.05.026>
- Boc A, Diallo AB, Makarenkov V (2012) T-rex: a web server for inferring, validating and visualizing phylogenetic trees and networks. *Nucleic Acids Res* 40(W1):W573
- Bollobás B (2001) *Random graphs*, Cambridge studies in advanced mathematics, 2nd edn. Cambridge University Press, Cambridge. <https://doi.org/10.1017/CBO9780511814068>
- Cadwalladr C (2017) The great British Brexit robbery: how our democracy was hijacked. *The Guardian*. The Observer. 2017, May 7. <https://www.theguardian.com/technology/2017/may/07/the-great-british-brexite-robbery-hijacked-democracy>. Accessed on 28 May 2017
- Chermak S, Freilich J, Suttmoeller M (2013) The organizational dynamics of far-right hate groups in the United States: comparing violent to non-violent organizations. *Stud Conflict Terror* 36:193–218
- Craig MA, Richeson JA (2014) On the precipice of a ‘majority–minority’ America: perceived status threat from racial demographic shift affects White Americans’ political ideology. *Psychol Sci* 25(6):1189–1197
- Crandall CS, White MHI (2016) Trump and the social psychology of prejudice. *Undark*. VARIABLES/Essays & Opinion, 11.17.2016. <https://undark.org/article/trump-social-psychology-prejudice-unleashed/>. Accessed 28 May 2017
- Dawkins R (1976) *The selfish gene*. Oxford University Press, Oxford
- Detrow S (2016) KKK paper endorses Trump; campaign calls outlet repulsive. *NPR*. November 2, 2016 10:05 AM ET. <http://www.npr.org/2016/11/02/500352353/kkk-paper-endorses-trump-campaign-calls-outlet-repulsive>. Accessed 28 May 2017
- Developers T (2018) Twitter developer documentation. <https://dev.twitter.com/docs>. Accessed 17 May 2017
- Dress AW (1984) Trees, tight extensions of metric spaces, and the cohomological dimension of certain groups: a note on combinatorial properties of metric spaces. *Adv Math* 53(3):321–402
- Dress A, Terhalle W (1996) The real tree. *Adv Math* 120(2):283–301
- Editors (2017) The Trump tweet tracker. *The Atlantic*. Live BLOG May 18, 2017. <https://www.theatlantic.com/liveblogs/2017/05/donald-trump-twitter/511619/>. Accessed 28 May 2017
- Enten H (2016) Americans’ distaste for both Trump and Clinton is record-breaking. *FiveThirtyEight.com*. May 5 2016 at 8:29 AM. <https://fivethirtyeight.com/features/americans-distaste-for-both-trump-and-clinton-is-record-breaking/>. Accessed 28 May 2017
- Ferrara E, Varol O, Davis C, Menczer F, Flammini A (2016) The rise of social bots. *Commun ACM* 59(7):96–104
- Fortuna P, Nunes S (2018) A survey on automatic detection of hate speech in text. *ACM Comput Surv* 51(4):85:1–85:30. <https://doi.org/10.1145/3232676>
- Freilich JDAPW (2006) Mismeasuring militias: limitations of advocacy group data and of state-level studies of paramilitary groups. *Justice Q* 23(1):147–162
- Gemignani M, Hernandez-Albujar Y (2015) Hate groups targeting unauthorized immigrants in the US: discourses, narratives, and subjectivation practices on their websites. *Ethn Racial Stud* 38(15):2754–2770
- Gilliard-Matthews S (2011) The impact of the economic downturn, immigrants, and political representation of White supremacist organization in the United States. *Sociol Focus* 44(3):255–279
- Glaser J, Dixit J, Green DP (2002) Studying hate crime with the internet: what makes racists advocate racial violence? *J Soc Issues* 58(1):177–193
- Graham D (2016) How did Maajid Nawaz end up on a list of Anti Muslim extremists? *The Atlantic*. October 29, 2016. <https://www.theatlantic.com/international/archive/2016/10/maajid-nawaz-splc-anti-muslim-extremist/505685/>. Accessed 28 June 2017
- Grassegger H, Krogerus M (2017) The data that turned the world upside down. *Stanford Public Policy Program*. VARIABLES/ Essays & Opinion, 01.28.2017. <https://publicpolicy.stanford.edu/news/data-turned-world-upside-down>. Accessed 15 June 2017
- Green DP, Rich A (1998) White supremacist activity and cross burning in North Carolina. *J Quant Criminol* 14:263–282
- Hale WC (2012) Extremism on the world wide web: a research review. *Crim Justice Stud* 25(4):343–356
- Hananoki E (2016) The complete history of Donald Trump’s relationship with the White nationalist movement. *Media Matters for America*. Blog, August 18, 2016 2:40 PM EDT. <https://www.mediamatters.org/blog/2016/08/18/complete-history-donald-trumps-relationship-white-nationalist-movement/212502>. Accessed 28 May 2017
- Hayden ME (2017) US civil rights groups worry anti-Muslim sentiment is fueling right-wing extremism. *ABC News*. Jun 12, 2017, 6:57 AM ET. <http://abcnews.go.com/US/us-civil-rights-groups-worry-anti-muslim-sentiment/story?id=47840271>. Accessed 15 June 2017

- Hochschild AR (2016) Strangers in their own land: anger and mourning on the American Right. New Press, New York
- Hogg MA (2007) Uncertainty–identity theory. *Adv Exp Soc Psychol* 39:69–126
- Hogg MA, Reid SA (2006) Social identity, self-categorization and communication of group norms. *Commun Theory* 16(1):7–30
- Howard PN, Kollanyi B, Woolley S (2016) Bots and automation over Twitter during the US election. Zugriff am 14. <http://comprop.oii.ox.ac.uk/wp-content/uploads/sites/89/2016/11/Data-Memo-US-Election.pdf>
- Ingraham C (2016) Two new studies find racial anxiety is the biggest driver of support for Trump. *Washington Post*. Wonkblog, June 6 2016. https://www.washingtonpost.com/news/wonk/wp/2016/06/06/racial-anxiety-is-a-huge-driver-of-support-for-donald-trump-two-new-studies-find/?utm_term=.d5b2d1ca159b. Accessed 28 May 2017
- Jansen BJ, Zhang M, Sobel K, Chowdury A (2009) Twitter power: tweets as electronic word of mouth. *J Am Soc Inf Sci Technol* 60(11):2169–2188
- Jonsson P (2013) Annual report cites rise in hate groups, but some ask: what is hate? *Christian Science Monitor*. February 23, 2011. <http://www.csmonitor.com/USA/Society/2011/0223/Annual-report-cites-rise-in-hate-groups-but-some-ask-What-is-hate>. Accessed 28 May 2017
- Jost JT, Glaser J, Sulloway FJ (2003) Political conservatism as motivated social cognition. *Psychol Bull* 129(3):339–375
- Kharakh B, Primack D (2016) Donald Trump's social media ties to White supremacists. *Fortune*. March 22, 2016, 10:03 AM EDT. <http://fortune.com/donald-trump-white-supremacist-genocide/>. Accessed 28 May 2017
- Knowles ED, Tropp LR (2016) Donald Trump at the rise of White identity politics. *Huffington Post*. THE BLOG 10/21/2016 07:15 am ET—Updated Oct 21, 2016. http://www.huffingtonpost.com/the-conversation-us/donald-trump-and-the-rise_b_12584616.html. Accessed 28 May 2017
- Kramer S (2017) Identifying viral bots and cyborgs in social media. Tech. rep. <https://www.oreilly.com/ideas/identifying-viral-bots-and-cyborgs-in-social-media>. Accessed 18 Feb 2018
- Kwak H, Lee C, Park H, Moon S (2010) What is Twitter, a social network or a news media? In: *Proceedings of the 19th international conference on world wide web*, WWW 10, pp 591–600. ACM, New York
- Leskovec J, Horvitz E (2008) Planetary-scale views on a large instant-messaging network. In: Huai J, Chen R, Hon HW, Liu Y, Ma WY, Tomkins A, Zhang X (eds) *WWW*. ACM, pp 915–924
- Major B, Blodorn A, Major Blascovich G (2017) The threat of increasing diversity: why many White Americans support Trump in the 2016 presidential election. *Group Process Intergroup Relat* 21:931–940 (in press)
- Martin T (2017) Dissecting Trump's most rabid online following. *FiveThirtyEight*. <https://fivethirtyeight.com/features/dissecting-trumps-most-rabid-online-following/>. Accessed 14 June 2017
- Metaxas P, Mustafaraj E, Wong K, Zeng L, O'Keefe M, Finn S (2015) What do retweets indicate? Results from user survey and meta-review of research. In: *International AAAI conference on web and social media*. ACM
- Milgram S (1967) The small-world problem. *Psychol Today* 1(1):60–67
- Mondal M, Silva LA, Benevenuto F (2017) A measurement study of hate speech in social media. In: *Proceedings of the 28th ACM conference on hypertext and social media*, HT '17. ACM, New York, pp 85–94. <https://doi.org/10.1145/3078714.3078723>
- Müller K, Schwarz C (2018) Making america hate again? Twitter and hate crime under Trump. Tech. rep. <https://doi.org/10.2139/ssrn.3149103>. Accessed on 20 May 2019
- Newman MEJ, Strogatz SH, Watts DJ (2001) Random graphs with arbitrary degree distributions and their applications. *Phys Rev E* 64:026118
- O'Connor T (2017) Trump, alt-right movement: Nazis, white supremacists praise new security measures against radical Islam. *International Business Times*. <https://tinyurl.com/y9j8oq72>. Accessed 21 July 2017
- Packer G (2016) Head of the class. *The New Yorker*. COMMENT May 16 2016 issue. <http://www.newyorker.com/magazine/2016/05/16/how-donald-trump-appeals-to-the-white-working-class>. Accessed 28 May 2017
- Pew Research Center (2015) Beyond distrust: how Americans view their government. Research report, Pew Research Center. <http://www.people-press.org/2015/11/23/beyond-distrust-how-americans-view-their-government/>. Accessed 28 May 2017
- Piggott S (2017) Neo-Nazi Billy Roper Organizing ACT for America's "March Against Sharia" in Arkansas. *Southern Poverty Law Center*. Hate Watch, June 6, 2017. <https://www.splcenter.org/hatewatch/2017/06/06/neo-nazi-billy-roper-organizing-act-america-march-against-sharia-arkansas>. Accessed 15 June 2017
- Raymond L, Pyke A (2017) "Trump is setting us free": white supremacists celebrate reports that Trump will dial down scrutiny. *Think Progress*. <https://tinyurl.com/znmy9xe>. Accessed 28 May 2017
- Reicher SD, Haslam SA (2017) Trump's appeal: what psychology tells us. *Scientific American Mind*, New York
- Reuters (2016) Ku Klux Klan newspaper declares support for Trump. *Reuters*. POLITICS—Wed Nov 2, 2016—4:22 pm EDT, Reporting by Mohammad Zargham; Editing by Jonathan Oatis. <http://www.reuters.com/article/us-usa-election-trump-kkk-idUSKBN12X2IG>. Accessed 28 May 2017
- Roussos G (2016) Donald Trump's psychological manipulation of the American people. Research blog, Yale Institution for Social and Policy Studies. <https://tinyurl.com/y9yacnyr>. Accessed 28 May 2017
- Runciman W (1966) Relative deprivation and social justice: a study of attitudes to social inequality in twentieth-century England. Reports of the Institute of Community Studies. University of California Press
- Ryan ME, Leeson PT (2011) Hate groups and hate crimes. *Int Rev Law Econ* 31(4):256–262
- Sainudiin R, Welch D (2016) The transmission process: a combinatorial stochastic process for the evolution of transmission trees over networks. *J Theor Biol* 410:137–170
- Sainudiin R, Yogeeswaran K, Nash K, Sahioun R (2018) Rejecting the null hypothesis of apathetic retweeting of US politicians and SPLC-defined hate groups in the 2016 US presidential election. In: Brandes U, Reddy C, Tagarelli A (eds) *IEEE/ACM 2018 international conference on advances in social networks analysis and mining*, ASONAM 2018, Barcelona, Spain, August 28–31, 2018. IEEE Computer Society, pp 250–253. <https://doi.org/10.1109/ASONAM.2018.8508555>
- Saitou N, Nei M (1987) The neighbor-joining method: a new method for reconstructing phylogenetic trees. *Mol Biol Evol* 4(4):406
- Schmidt A, Wiegand M (2017) A survey on hate speech detection using natural language processing. In: *Proceedings of the fifth international workshop on natural language processing for social media*. Association for Computational Linguistics, Valencia, pp 1–10. <https://doi.org/10.18653/v1/W17-1101>. <https://www.aclweb.org/anthology/W17-1101>
- Selk A (2017) Twitter co-founder: I'm sorry if we made Trump's presidency possible. *Washington Post*. May 21, 2017. <https://tinyurl.com/yckypjku>. Accessed on 28 May 2017
- Shalby C (2016) A brief history of the Trump campaign's controversies with women. *Los Angeles Times*. October 9, 2016, 7:00 AM. <http://www.latimes.com/politics/la-na-pol-trump-campaign-insults-women-20161009-snap-htmlstory.html>. Accessed 28 May 2017

- Solon O (2013) Richard Dawkins on the internet's hijacking of the word 'meme'. WIRED Magazine. 20 June 13. <http://www.wired.co.uk/news/archive/2013-06/20/richard-dawkins-memes>. Accessed 22 June 2017
- Southern Poverty Law Center (2016) Hate map. SPLC. October 11, 2013 4:00 AM. <https://www.splcenter.org/hate-map>. Accessed 28 May 2017
- Southern Poverty Law Center (2017) The year in hate and extremism. SPLC. <https://www.splcenter.org/fighting-hate/intelligence-report/2017/year-hate-and-extremism>. Accessed 21 May 2019
- Stephan WG, Stephan CW (2000) An integrated threat theory of prejudice. In: Oskamp S (ed) Reducing prejudice and discrimination, Claremont symposium on applied social psychology series, chapter 2. Taylor & Francis, pp 23–45
- Torres A (2017) To the Southern Poverty Law Center, believing homosexual activity is sinful is 'hate' and 'bigotry'. National Review (2013). October 11, 2013. <http://www.nationalreview.com/article/360968/splc-and-hate-groups-alec-torres>. Accessed 28 May 2017
- Tumasjan A, Sprenger T, Sandner P, Welp I (2010) Predicting elections with Twitter: what 140 characters reveal about political sentiment. In: Proceedings of the fourth international AAAI conference on weblogs and social media, pp 178–185
- Tuttle I (2017) Racism (still) didn't elect Trump. National Review. <http://www.nationalreview.com/corner/446808/washington-post-racism-motivated-trump-voters>. Accessed 14 June 2017
- Twitter Public Policy (2018) update on Twitter's review of the 2016 US election. Tech. rep. https://blog.twitter.com/official/en_us/topics/company/2018/2016-election-update.html. Accessed 18 Feb 2018
- Vaccari C, Valeriani A, Barbera P, Bonneau R, Jost JT, Nagler J, Tucker JA (2015) Political expression and action on social media: exploring the relationship between lower and higher-threshold political activities among Twitter users in Italy. *J Comput Mediat Commun* 20(2):221–239
- Walker I, Pettigrew TF (1984) Relative deprivation theory: an overview and conceptual critique. *Br J Soc Psychol* 23:301–310
- Watts DJ, Strogatz SH (1998) Collective dynamics of 'small-world' networks. *Nature* 393:440–442
- Webster DM, Kruglanski AW (1994) Individual differences in need for cognitive closure. *J Personal Soc Psychol* 67(6):1049–1062
- Wood T (2017) Racism motivated Trump voters more than authoritarianism. Washington Post. <https://tinyurl.com/yavzdm5s>. Accessed 14 June 2017
- Woolley S, Guibeault D (2017) Computational propaganda in the United States of America: manufacturing consensus online. Samuel Woolley and Philip Howard, Eds. Working Paper No. 2017.5 Project on Computational Propaganda, Oxford, UK, pp 1–29. <http://comprop.oii.ox.ac.uk/wp-content/uploads/sites/89/2017/06/Comprop-USA.pdf>
- Yogeeswaran K, Dasgupta N (2014) The devil is in the details: abstract versus concrete construals of multiculturalism differentially impact intergroup relations. *J Personal Soc Psychol* 106(5):772–789
- Zaharia M, Xin RS, Wendell P, Das T, Armbrust M, Dave A, Meng X, Rosen J, Venkataraman S, Franklin MJ, Ghodsi A, Gonzalez J, Shenker S, Stoica I (2016) Apache spark: a unified engine for big data processing. *Commun ACM* 59(11):56–65

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