

Inside the Right-Leaning Echo Chambers: Characterizing Gab, an Unmoderated Social System

Lucas Lima*, Julio C. S. Reis*, Philippe Melo*, Fabricio Murai*,
Leandro Araújo*, Pantelis Vikatos[†], Fabrício Benevenuto*

*Universidade Federal de Minas Gerais (UFMG), Brazil

[†]IIT-CNR, Pisa, Italy

{lucaslima, julio.reis, philipe, murai, leandroaraujo, fabricio}@dcc.ufmg.br, p.vikatos@iit.cnr.it

Abstract—The moderation of content in many social media systems, such as Twitter and Facebook, motivated the emergence of a new social network system that promotes free speech, named Gab. Soon after that, Gab has been removed from Google Play Store for violating the company’s hate speech policy and it has been rejected by Apple for similar reasons. In this paper we characterize Gab, aiming at understanding who are the users who joined it and what kind of content they share in this system. Our findings show that Gab is a very politically oriented system that hosts banned users from other social networks, some of them due to possible cases of hate speech and association with extremism. We provide the first measurement of news dissemination inside a right-leaning echo chamber, investigating a social media where readers are rarely exposed to content that cuts across ideological lines, but rather are fed with content that reinforces their current political or social views.

Index Terms—Gab, Social Media, Echo Chambers, News Sharing, Online News

I. INTRODUCTION

Free speech and hate speech have opened a long debate between regulation and freedom of expression, especially in the online space. Currently, most online social networks have strict policies against hate speech and moderate the content posted by its users. In particular, Twitter, Facebook, Google (YouTube), and Reddit have largely increased removals of hate speech [1, 2]. This scenario has motivated the emergence of a new social network system, named Gab¹, which was originally designed to be a social network that promotes free speech.

In spite of advocating liberty and freedom of speech, Gab has received several criticisms regarding the content shared there. The lack of moderation in this system started to attract users banned due to hate speech from other social networks [3]. Gab has also been pointed as a safe space for alt-right groups [4]. This environment favors the emergence of clusters of like-minded individuals and the polarization of opinions, creating information bubbles around users who are fed with information that reinforces their current political or social views. This phenomenon is known as the “echo chamber effect” and it has been extensively studied recently [5, 6].

Several efforts have proposed ideas to mitigate the effect of “echo chambers” or “filter bubbles”, either by introducing diversity in the news that users are consuming [7, 8] or by

highlighting posts that evoke similar reactions from opposite political views [9]. Despite the importance of all these efforts, little is still known about exactly what happens inside an echo chamber. This lack of understanding is likely due to the difficulty of splitting apart what is a bubble in social networks such as Facebook and Twitter.

Therefore, understanding what happens inside Gab can provide us valuable insights on the study of right-leaning echo chambers. Our effort consists of answering the following research questions concerning Gab users and posts:

Research Question 1 (RQ1): Who are the users who joined Gab in terms of extremist views, political leaning and ability to spread news?

Research Question 2 (RQ2): What kind of content is shared in Gab? What are the news shared within this system?

Our findings show that Gab is a very politically oriented system, in agreement with concurrent work [10]. Furthermore, we show that the majority of Gab users are conservative, male, and Caucasian. We were able to identify many users listed as extremists by the main media who showed to be influential and very active in the Gab network. The most popular type of discussion in Gab posts is focused on politics and conservatism. We also identified a few cases in which messages are considered highly toxic.

We also show that, although most of the (unique) news domains present in shared URLs are considered left-leaning, conservative news outlets comprise a larger fraction of the shared links. By quantifying the popularity of sites whose links are disseminated inside this ideology-biased community, we show that popular news domains shared on Gab are not popular on Facebook or on the Internet.

The rest of the paper is organized as follows. Section II briefly reviews related work. Then, we present our datasets and the strategies we used to collect them in Section III. Section IV and Section V provide an in-depth investigation of our two research questions. Section VI provides our conclusions.

II. RELATED WORK

In this section, we review related work along two dimensions: (i) studies related to freedom of expression and (ii) efforts related to news sharing and propagation in social media.

¹<https://gab.ai/>

A. Freedom of Expression

Freedom of speech is guaranteed by the first amendment in the USA and by other similar laws in liberal countries. A key argument pointed out by supporters of absolute free speech is that it is impossible to give power to anyone to restrict hate speech without allowing bans against non-hate speech. On the other hand, other theories defend that even free speech needs regulation. For instance, back in 1945, Karl Popper coined the paradox of tolerance [11] as *Unlimited tolerance must lead to the disappearance of tolerance*. The key insight from these arguments is that if a society is tolerant without limit, their ability to be tolerant will eventually be seized or destroyed by the intolerants. More recently, Mitchell [12] analyzes how freedom to speak is regulated in liberal states and points liberal states guarantees freedoms, which require the creation of legal controls and regulations to ensure the social order.

These paradoxes and arguments, however, are still just theories, as there has been a lack of empirical data to support them. More important, many social systems have recently seen themselves in the midst of these paradoxes. As hate speech grows in the Web [13, 14], they needed to create hate speech policies and regulate what is expressed in their systems. The data within Gab now allow researchers to empirically measure what really happens in an unlimited free speech environment. Thus, by characterizing this system, we hope to bring insights that can contribute to this discussion.

B. News Sharing and Propagation

The dissemination of news in social networks has been the subject of several studies [15] focusing on different aspects, such as bias [16] and political news [17], or yet, the characteristics of players (or spreaders) [18]. In the case of Twitter, for example, Rodrigues *et al.* [19] show that retweets are responsible for increasing the audience of URLs by about two orders of magnitude.

The characteristics (or factors) influencing news sharing in social media are explored by Lee and Ma [20]. Based on uses and gratifications, social cognitive theories, and interviews, the authors explore the influence of information seeking, socializing, entertainment, status seeking and prior social media sharing experience on news sharing intention. More recently, Garimella *et al.* [21] study political echo chambers on social media. They examine the interaction between shared opinions and dissemination, dividing Twitter users in three categories: (i) users who are exposed to content that reinforces their opinions, (ii) users who try to bridge echo chambers by sharing content with diverse leaning, and (iii) gatekeepers, which are users who consume diverse content but produce content that fuels the bubbles. Differently, Reis *et al.* [22] investigate the demographics of users in Twitter and how it affects news sharing. They show that white and male users tend to share more news on Twitter, biasing the news audience towards the interests of these demographic groups.

Finally, in a very recent work, Ottoni *et al.* [23] investigate the presence of hateful content and discriminatory bias in

right-wing YouTube channels, contributing to a better understanding of the behavior of right-wing YouTube users.

Overall, our study is complementary to previous efforts as our work provides an in-depth investigation of what kind of news are shared in a right-leaning echo chambers, thus providing a broader understanding of Gab users and content.

III. METHODOLOGY

In order to characterize Gab as a right-leaning echo chamber and to analyze its news sharing ecosystem, we describe next our strategy for data collection. We start by presenting an overview of the social network.

A. The Gab Social Network

In essence, Gab is very similar to Twitter, but barely moderates any of the content shared by its users. According to Gab guidelines, the site does not allow illegal pornography and promotion of violence and terrorism². Everything else towards free speech is allowed. Posts are limited to 3,000 characters and there is an upvoting and downvoting system for posts, which does not exist on Twitter. Additionally, posts are categorized as News, Politics, Art etc., that users can use to find popular content on specific topics.

B. Data Collection

Our dataset includes data from users profiles, their lists of followers and friends, and all their messages. Each user in our dataset has an user ID, name, screen name, number of friends, followers, posts, account creation date, information of whether the profile is verified, PRO (subscribers), or Premium (users paid for their content), profile picture URL, and a profile bio. Each post comes along with a timestamp and its category.

We crawled users and their posts by following a Breadth-First Search (BFS) scheme on the graph of followers and friends. We used as seeds users who authored posts listed by categories in the Gab main page. We implemented a distributed and parallel crawler that ran in August 2017 which took only three days to gather all users reachable from the seeds. In total, our dataset comprises **171,920** users (the estimated number of users in August 2017 was 225 thousand [24]) and a total of **12,829,976** posts.

After collected, we preprocessed these data by discarding all messages marked as repost, quotes or unlabeled messages (i.e., messages not associated with any category). After these steps, we reduced the data to a subset containing all the relevant information to our characterization: a sample of **1,297,165** posts incorporating useful information for analyses discussed in the following sessions.

News Dataset. In order to provide a measurement of news dissemination inside Gab, we collected URLs' headlines of posts categorized as news containing URLs (315, 935) according to the following two steps. First, we attempt to parse the HTMLs retrieved by the links and search for specific tags, such as `<title>`. When this is not possible, we match regular expressions to retrieve potential headlines of links from its own

²<https://gab.ai/about/guidelines>

URL. Following these steps, 76.8% out of the total of news messages have an associated headline. The following sections delve deeper in the study of these URLs, its domains, and headlines.

IV. RQ1: WHO JOINED GAB?

In this section, we provide a series of analyses aiming at depicting who are the users that joined Gab.

A. When did users join Gab?

Since the release of the service in August 2016, Gab has experienced a constant and slow increase in its number of users, except for a large peak around November 2016, during the US presidential election period. From all users we crawled, 51.4% joined Gab from October to December, in 2016, 28.7% of them just in November 2016. This suggests that the US presidential election represented an important motivation for users to join Gab. Zannettou *et al.* [10] also highlight that Gab has become popular among users in the occurrence of political events such as the 2017 Charlottesville Unite the Right rally.

Next, we attempt to infer the ideological leaning of Gab users.

B. Political Leaning of users on Gab

In order to measure the political bias of Gab users, we used a recently introduced framework [25], kindly shared by the authors. Their approach measures the political leaning of Twitter users, given a set of information from their friends. Thus, to identify the political leaning of Gab users, we first developed a strategy to identify the Twitter account of Gab users.

1) *Identifying Gab users on Twitter:* We design a two-step methodology. First, we searched for the screen name of all Gab users in the *Twitter REST API*³, identifying 62,291 (36.23%) users with exactly the same screen name in both systems. Then, we compare the profile name of these users from Gab with their respective profile names in Twitter, keeping only those for which we found an exact match. Our final dataset contains 23,030 users, which corresponds to about 13% of the total number of users.

2) *Inferring Political Bias of Twitter users:* For each user identified as being the same on Twitter and Gab, we further gathered her lists of Twitter friends from the Twitter API. Then, the framework used to measure these users' political leaning [25] was able to identify the leaning of 16,804 users (73% of the matched pairs).

Figure 1 shows the distribution of bias scores. The closer the bias score is to +1, the more liberal a user is. Users with bias score between -0.03 and +0.03 are considered neutral or moderate. We note a large number of extremely conservative users (i.e. those with scores close to -1) in Gab and most of the liberal ones have ideological leaning scores quite close to 0 (i.e. moderate). Out of the 16,804 users we sampled, 6,237 (37.1%) were inferred to be conservative, 2,925 (17.4%) as liberal, and 7,642 (45.5%) as moderate.




Fig. 1: Distribution of bias scores for mapped users.

Hence, the ratio between conservatives and liberals is about 2.13.

In order to contrast these values with those from other social networks, we use the Facebook audience API, following the same methodology of recent efforts [26]. Facebook provides this API for advertisers to estimate the amount of users who are likely to match an advertising criteria⁴. We then used the Audience API to gather the political leaning of all Facebook users in the US. Overall, the inferred fraction of conservatives in Facebook is 33.6%, 42.6% of liberals and 23.8% of moderates. The ratio between conservatives and liberals is about 0.79. Therefore, in comparison with Facebook, Gab has a much more conservative than liberal crowd.

One limitation of this analysis is that it requires us to match Gab and Twitter accounts as a first step, thus reducing the set of suitable users to 13.4%. This is due to the fact the framework provided by Kulshrestha *et al.* [25] is a Twitter-based approach. We leave as future work the creation of a new inference framework based on the Gab network only.

C. Gender and Race of Gab users

In order to better understand demographic aspects of users in Gab, we used a methodology recently explored in previous efforts [27, 28] to infer the gender and race of Gab users. The strategy consists of gathering the profile picture Web link of Gab users and submit these links into the Face++ API. Face++ is a face recognition platform based on deep learning which is able to identify gender (i.e. male and female) and race (limited to Asian, Black, and White) from recognized faces in images. We have discarded users whose profile does not have a face picture or does not have a recognizable face according to Face++. A recent work [28] evaluated the effectiveness of the inference made by Face++, using human annotators to label randomly selected profile images from Twitter. They measured the inter-annotator agreement in terms of the Fleiss' K score which was 1.0 and 0.865 for gender and race, respectively.

From the total number of Gab users we crawled, less than a half (47.22%) have a profile picture Web link. Given these

³<https://developer.twitter.com/en/docs>

⁴developers.facebook.com/docs/marketing-api/audiences-api

Race	Gender		Total
	Male	Female	
Asian	3,676 (10.4%)	1,920 (5.4%)	5,596 (15.8%)
Black	2,106 (5.9%)	787 (2.2%)	2,893 (8.2%)
White	18,078 (50.9%)	8,926 (25.1%)	27,004 (76.1%)
Total	23,860 (67.2%)	11,633 (32.7%)	35,493 (100%)

TABLE I: Demographic distribution of nearly 36 thousand Gab users.

82,215 users, only 35,493 have valid profile pictures that we were able to collect demographic information. Table I reports the demographic distribution of the 35,493 users in our dataset across the different demographic groups. We observe a prevalence of males (67.2%) in comparison to females (32.7%) and a high predominance of Whites (76.1%) in comparison to Asians (15.8%) and Blacks (8.2%). This means that if we pick users randomly in our dataset, we would expect demographic groups with these proportions. These proportions, in particular the ones regarding race, differ from Facebook, where the White men correspond to 29.35% of the total population [26].

D. Extremism in Gab

Next, we want to analyze whether far-right groups are present in Gab. According to the SPLC⁵, which is an American nonprofit legal advocacy organization specializing in civil rights dedicated to fighting hate, extremists in the United States come in different forms and follow a wide range of ideologies. As we look through the individual profiles of extremists provided by the SPLC [29] and the Anti-Defamation League (ADL) [30], we notice the presence of a few members of these movements in the platform. Table II presents the list of 29 listed extremist users by the two sources with Gab accounts and their number of Gab followers as observed in August 2017. In particular, most people listed by the ADL (61.11%) have a Gab account, reinforcing that this social network has been attracting a large number of extremely conservative users.

A key question that arises about these listed extremist users in Gab is whether they are widely followed and whether they are active users in the system. Among the top 10 most followed users, four are extremist users according to the SPLC and ADL, and all of them have verified accounts (that is a form of verifying identities). These considered extremist users, namely Milo Yiannopoulos (@m), Mike Cernovich (@Cernovich), Brittany Pettibone (@BrittPettibone) and Vox Day (@voxday), have on average 27,309.75 Gab followers. The average number of followers over the list of 29 users is 6,160.2, two orders of magnitude larger than the average number of followers for all Gab users we crawled (72.23). This means that posts of listed extremist users can reach a large number of different users within one hop. Approximately 35% of Gab users we crawled follow at least one of these 29 extremist users.

In regards to the number of posts, we also note that these users share in average more posts than the average Gab user. The average number of Gab posts over all users is 55.69,

Screen name	User name	# of followers
RealAlexJones	Alex Jones**	14,962
Alex_Linder	Alex Linder	963
AndrewAnglin	Andrew Anglin**	4,853
AndyNowicki	Andy Nowicki	75
thewizardofthorntonpark	Augustus Invictus	128
BillyRoper	Billy Roper	407
occdissent	Brad Griffin	534
BrittPettibone	Brittany Pettibone*	23,335
Cantwell	Christopher Cantwell***	2,728
DanielFriberg	Daniel Friberg	299
RealDavidDuke	David Duke****	982
GavinMcInnes	Gavin McInnes	2,681
Posobiec	Jack Posobiec*	2,944
jartaylor	Jared Taylor*	147
TheMadDimension	Jason Kessler*	454
mattforney	Matt Forney**	4,322
matthewheimbach	Matthew Heimbach	43
mattparrott	Matthew Parrott	56
Cernovich	Mike Cernovich*	27,462
mikeenoch	Mike Peinovich	703
m	Milo Yiannopoulos*	39,891
Pamela	Pamela Geller*	3,238
ramzpaul	Paul Ray Ramsey	2,951
pax	Pax Dickinson	12,218
Richardbspencer	Richard B Spencer	6,833
RobertSpencer	Robert Spencer	1,317
TaraMcCarthy	Tara McCarthy***	5,565
voxday	Theodore Beale*	18,551
Microchimp	Tim Gionet	4

* verified; ** verified and PRO; *** verified, PRO and premium; **** PRO

TABLE II: List of profiles considered extremist by the SPLC and ADL who were found on Gab.

whereas the average for listed extremist users is 571.66. This indicates that these users are not only more followed, but they are also more active in the system.

E. News spreaders

We now analyze the profile of news spreaders in our dataset, i.e. people who share news sources within the news category in Gab posts. It is interesting to note that the majority of the active users within the news category (60.14%) have posted at least one URL and, from these, 62.71% have posted more than one URL.

Table III shows the top 10 news spreaders and their total number of posts as of August 2017. These users have shared on average 10,838.5 posts. The user *Constitutional Drunk*, associated with the right-biased USSA News website, had the largest number of posts, 59,378. In fact, most publications from this user share content of that website, which happens to be the most shared domain. Only 37.9% of the users listed as extremists have posts categorized as news.

Aiming at understanding the influence of these users in Gab, we also analyze their number of followers. While the average number of followers among all crawled users is low (72.23), the top 10 news spreaders have on average 4,128.8 followers, two orders of magnitude higher than the overall mean. In total, 20,176 out of 171,920 Gab users follow at least one of these news spreaders, which implies that posts from the top 10 news spreaders can reach 11.74% of the the social network users.

⁵Southern Poverty Law Center <https://www.splcenter.org/>

Screen name	User name	# of posts
USSANews	Constitutional Drunk	59,378
Zlatford	Zak	13,388
wrathOfkhan	wrath Ofkhan	8,189
histanvan	Harry2	6,459
Kek_Magician	Kek_Magician	4,248
Lakeem	Lakeem Khodra***	4,016
Arwen777	Dani	3,913
weeklyflyer	Jerry	3,213
rabite	Stankpipe	2,925
OpenQuotes	OpenQuotes	2,656

*** verified, PRO and premium.

TABLE III: Top 10 news spreaders.



Fig. 2: Word clouds for Gab posts.

V. RQ2: WHAT USERS SHARE IN GAB?

In this section, we analyze the content shared in Gab. We start our analysis by characterizing typical language usage.

A. Popular Words and Topics in Gab

Figure 2 illustrates through one word cloud⁶ frequent terms of Gab posts for all posts from our dataset. Overall, this word cloud shows a strong political emphasis in Gab posts, with the existence of several terms which have been used by conservative political campaigns and their supporters. In particular, the hashtags *MAGA*, which stands for *Make America Great Again*, and *DRAIN THE SWAMP* were popularly mentioned as part of the Donald Trump’s 2016 presidential campaign.

The analysis of posts’ categories reinforces the strong political trait of the Gab social network. The category with the largest number of posts is News (35.74%), usually associated with politics, followed by the categories Politics, Humor, AMA (Ask Me Anything) and Entertainment. The other categories – Music, Technology, Art, Sports, Faith, Philosophy, Photography, Science, Finance and Cuisine – sum up to 225,611 (17.39%) posts.

B. Evidence of Hate Speech in Gab

Recently, there have been several efforts towards identifying online hate speech, as the problem manifests itself in different systems and under different forms. A recent API called Perspective⁷, created by Jigsaw and Google’s Counter Abuse Technology team, uses machine learning models to identify whether comments could be perceived as toxic or not,

⁶<http://www.wordle.net/>

⁷<https://www.perspectiveapi.com/#/>




Fig. 3: Cumulative distribution function (CDF) for toxicity scores of categorized posts.

providing toxicity scores which ranges from 0 to 1. A toxic comment is defined by the creators of the project as *a rude, disrespectful, or unreasonable comment that is likely to make you leave a discussion*. Given that properly detecting all the hate speech in Gab represents a methodological challenge as current methods for that are still very incipient, we provide evidence of hate speech in Gab through the analysis of Gab categorized posts’ toxic scores.

Figure 3 shows the cumulative distribution function (CDF) for the sample of 982,875 Gab posts (75.77% of the total number of categorized posts) from which we were able to gather toxicity scores from the API. Toxicity scores lower than 0.4 indicates posts unlikely to be perceived as toxic, while scores higher than 0.7 indicates posts likely to be perceived as toxic. Toxicity scores between 0.4 and 0.7 indicates unsureness whether posts will be perceived as toxic. We notice that more than 90% of these posts (91.4%) have toxicity scores less than 0.70, i.e. a large fraction of the total number of posts are either unlikely to be perceived as toxic, or may or may not be perceived as toxic. There are only 23,188 posts (2.4%) with toxicity score higher than 0.9.

Even though there is just a small percentage of Gab posts perceived as toxic, some posts with scores closer to 1 clearly incorporate hate speech towards groups of people of different religion and ethnic origin, race, sexual orientation and gender, e.g., *I F*CKING⁸ HATE N*GGERS THEY ARE ALL DISGUSTING AND THEY SHOULD ALL DIE A BUNCH OF F*CKING C*NTS* and *F*ck Islam! Muhammad was a f*ggot pedophile! F*ck Islam! F*ck all of it! Your religion s*cks, your culture s*cks, your women s*ck, every f*cking thing about Islam is inferior, barbaric, and gay. F*CK ISLAM*. Therefore, we notice there are users who abuse the lack of moderation to disseminate hate.

⁸Wherever present, the “*” has been inserted by us, in order to lessen the impact of the offense words may inflict on some people, and was not part of the original text.

Dataset	Republican (%)	Democrat (%)	Neutral (%)
Budak <i>et al.</i>	4 of 4 (100.00)	11 of 11 (100.00)	-
Mitchell <i>et al.</i>	6 of 7 (85.71)	19 of 23 (82.61)	1 of 3 (33.33)
Bakshy <i>et al.</i>	165 of 205 (80.49)	203 of 260 (78.08)	-
AllSides	29 of 40 (72.50)	26 of 40 (65.00)	24 of 32 (75.00)

TABLE IV: Number of domains found in Gab posts which are categorized as news and coexistence in each dataset.

C. What kind of news are shared in Gab?

Inferring bias of news sources often relies on strategies which combine analyses of the audience of news outlets and inspection of the published content. Recent research works such as Budak *et al.* [31], Mitchell *et al.* [32], Bakshy *et al.* [33], and AllSides⁹ infer the political leaning of over 600 news outlets using these techniques. Next, we briefly discuss each of these studies and, based on them, we analyze the presence of biased news sources on Gab.

Following a survey-based approach, Mitchell *et al.* [32] classify the audience of popular news media outlets based on a ten question survey covering a range of issues like homosexuality, immigration, economic policy, and the role of government, classifying the political leaning of the audience in five categories: consistently liberal, mostly liberal, mixed, mostly conservative, and consistently conservative. On the other hand, Bakshy *et al.* [33] follow a news-based approach in which the authors derive the alignment score of media outlets by first identifying the political leaning of over 10 million Facebook users based on self-declarations, and then considering how users with different political leanings shared the stories published by these outlets.

Other studies include approaches based on (a) content and (b) crowdsourcing. The former refers to the research of Budak *et al.* [31] where the authors use a content-based approach to identify the slant of the top 13 U.S. news outlets and two popular political blogs. Regarding the latter, the website AllSides.com infers bias by encouraging its users to rate different news outlets in one of the five categories: left, lean left, center, lean right, and right.

Table IV presents, for each of the aforementioned datasets, the number of news outlets inferred as Republican, Democrat, and Neutral, which were also shared on Gab posts categorized as news. We notice that regardless of the bias of news outlets, most domains are in fact found in Gab posts (e.g., 100% of the news outlets with political leaning inferred by Budak *et al.* have been shared in Gab). However, we show next that news sources inferred as Republican are shared on average more often than those inferred as Democrat and Neutral.

Figure 4 shows box plots of the number of times news sources inferred as Republican, Democrat, and Neutral have been shared on Gab posts categorized as news. Clearly, news outlets biased towards conservative audience are consistently shared on average more than the others across all four methodologies. In particular, considering only news outlets found in the study of Budak *et al.* (Figure 4a), Gab posts comprise on

⁹<https://allsides.com/media-bias/media-bias-ratings>

Source	Domain	(%)
USSA News	ussanews.com	16.40
Youtube	youtube.com	15.99
Breitbart	breitbart.com	5.17
Zero Hedge	zerohedge.com	2.08
Infowars	infowars.com	1.91
The Gateway Pundit	thegatewaypundit.com	1.59
Daily Mail Online	dailymail.co.uk	1.52
RT News	rt.com	1.48
The Daily Caller	dailycaller.com	1.25
Tyranny News	fightagainst-tyranny.com	1.21
EXPRESS	express.co.uk	1.20
Fox News	foxnews.com	0.99
Truth Feed News	truthfeed.com	0.63
Sputnik News	sputniknews.com	0.63
ABC News	abcnews.go.com	0.57
Reuters	reuters.com	0.50
Washington Examiner	washingtonexaminer.com	0.49
Behoerdenstress Whistleblower	behoerdenstress.de	0.46
The Hill	thehill.com	0.46
Worldwide Weird News	worldwideweirdnews.com	0.45
The Daily Informer	thedailyinformer.net	0.44
WND	wnd.com	0.44
New York Post	nypost.com	0.41
The Washington Post	washingtonpost.com	0.34
The Last Refuge	theconservativetreehouse.com	0.34
The Telegraph	telegraph.co.uk	0.33
BBC News	bbc.com	0.33
The Washington Free Beacon	freebeacon.com	0.32
INDEPENDENT	independent.co.uk	0.30
The Sun	thesun.co.uk	0.30

TABLE V: Top 30 news sources in posts and their respective domain, percentage over all posts.

average 5,894 right-leaning news sources (median 2,116) and only 645.6 leftist news sources (median 538).

D. What are the most shared domains?

First, we extract the domains of all links we collected from 463,663 Gab posts categorized as news. Table V shows the proportion of the top 30 most frequently shared domains in this social media¹⁰. We observe that ussanews.com is the most frequently shared news domain, accounting for more than 16% of the news URLs shared on Gab, followed by youtube.com, which accounts for more than 15%. Although YouTube is not specifically designed for news, this social media is widely used for publishing news videos [34, 35]. Moreover, we note a substantial presence of alternative news outlets like breitbart.com, zerohedge.com, infowars.com and thegatewaypundit.com, which correspond to more than 5%, 2%, 1.9% and 1.5%, respectively. Finally, the newspaper dailymail.co.uk (1.5%) figures among the most shared domains, followed by rt.com (1.4%), dailycaller.com (1.2%) and fightagainst-tyranny.com (1.2%).

Some news domains widely shared in Gab are not very popular according to Alexa.com. For example, the most shared domain in Gab (ussanews.com) appears in a very low position in the Alexa rank (409,260). The same is true for others domains such as fightagainst-tyranny.com, truthfeed.com, behoerdenstress.de, worldwideweirdnews.com and thedailyinformer.net.

¹⁰We exclude from this ranking domains not belonging to news outlets (i.e. twitter.com, i.imgur.com and pbs.twimg.com)

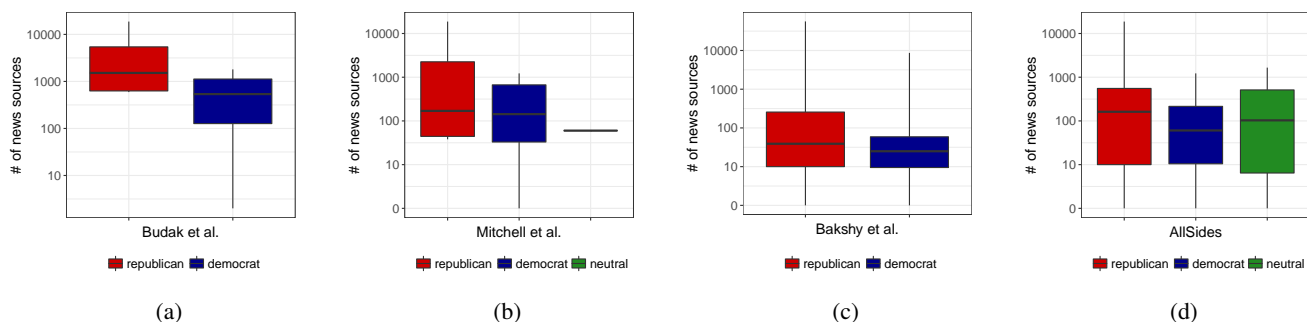


Fig. 4: Number of times news sources were shared in Gab posts categorized as news, grouped by political leaning (Republican, Democrat, and Neutral) as inferred by Budak et al. (a), Mitchell et al. (b), Bakshy et al. (c), and AllSides (d). Each box plot shows minimum, 25-percentile, median, 75-percentile and maximum.

URL News	Shares	Bit.ly
http://bit.ly/2qNboDV	99	949
https://youtu.be/4Emd7urcWBc	60	-
https://wikileaks.org/ciav7p1/	44	4,025
https://youtu.be/MHZSfhd1X_8	39	-
http://bit.ly/2EJ8WpZ	35	-
http://washex.am/2ivMBhv	35	772
http://dailym.ai/2kubLOX	32	0
http://bit.ly/2nZdrD7	27	-
http://dailym.ai/2qSZ5mi	25	-
http://bit.ly/2nB285x	24	4,492
http://bit.ly/2C0QdW1	24	1
http://bit.ly/2Es9CAT	24	-
http://bit.ly/2kSgBZC	24	7,356
http://bit.ly/2iH0mLy	24	274
http://bit.ly/2nZlXlK	24	-

TABLE VI: Top 15 most popular links shared in Gab, their number of shares in Gab and popularity according to Bit.ly (larger means more popular).

Overall, these results show that there is a large diversity of news domains shared in Gab. While most domains are shared very few times, a few domains comprise most of the shared URLs related to news. Interestingly, there is **no intersection** between the top 30 most shared domains in Gab and the top 10 most shared news domains in traditional media like Facebook¹¹. On Twitter, previous works showed that news URLs from online newspapers like “The New York Times” are shared very often [22]. However, this domain does not appear frequently in Gab shares, highlighting important differences between these social networks.

E. What are the top stories shared in Gab?

Next, we shift our focus to understand what are the top stories shared in the category news of the platform.

Table VI shows the 15 most popular links shared in Gab, i.e. the most frequently posted links. We also show their popularity according to Bit.ly¹², which we use as a proxy for the number of clicks made to each of these links. Bit.ly is a well-known URL shortening service that shortens millions of URLs daily. The service API provides the possibility of checking the total number of clicks that a shortened link has received [36]. In

¹¹<https://www.statista.com/statistics/265830/facebook-daily-newspapers-top-ten/>

¹²<http://bit.ly/>

general, these links presented in Table VI have been posted on average 36 times, where the first link in the rank appears in 99 posts and the last one appears in 24 posts. Some popular links shared in Gab have Bit.ly popularity higher than links to popular news sources such as BBC News, DailyMail, New York Times, and Reuters [37].

The most popular link is a petition created on May 19, 2017, requesting action of the Congress regarding the following issue: “Appoint a Special Prosecutor to investigate the murder of Seth Rich, the alleged Wikileaks email leaker”. Other popular links include YouTube videos (the second most popular link talks about the alternative media), WikiLeaks news, Donald Trump-related news, and also news which might not be completely accurate. For instance, the sixth most popular link has been rated by the fact-checking website Snopes as mostly false¹³. This suggests that Gab is prone to the dissemination of fake news, similarly to other social networks. However, the fake news that spread in Gab are usually those that reinforce right-leaning beliefs.

VI. CONCLUSION

In this paper we characterize Gab, a social network that emerged advocating liberty and freedom of expression, but received several criticisms about the content shared in it. As the debate between hate speech and free speech is an open and very controversial issue, Gab is an evident source of data for many researchers to explore deeper all those concerns. In this direction, our study provides a characterization of users and content shared on Gab and also contributes to understanding the behavior of such users without the strict policies of moderation from other medias.

Our findings show that Gab is a very politically oriented system that hosts known banned users from other social networks. We also show that the majority of Gab users are conservative, male, and Caucasian. Gab is also crowded by extremist users. Our analysis of what users share in Gab unveiled a lot of political statements. Posts indicate that, while users support free speech, a small part of the posts not only mirror political views but incorporate hate speech.

¹³<https://www.snopes.com/child-prostitution-legalized-in-california/>

Our analyses also show the existence of a great variety of news domains shared within Gab. Most of these domains are not popular in other social media or on the Internet as a whole, and part of them are known for spreading politics-related news. These results indicate that an unmoderated social media such as Gab has become an echo chamber for right-leaning content dissemination.

ACKNOWLEDGMENTS

This research was partly supported by CAPES, MASWeb (grant FAPEMIG/PRONEX APQ-01400-14), Fapemig, and Humboldt Foundation. We also thank Juhi Kulshrestha for measuring the political leaning of the users in Gab.

REFERENCES

- [1] A. Macdonald and J. Fioretti, "Social media firms have increased removals of online hate speech: EU," <https://www.reuters.com/article/us-eu-hatespeech/social-media-firms-have-increased-removals-of-online-hate-speech-eu-idUSKBN18S3FO>, May 2017.
- [2] E. Chandrasekharan, U. Pavalanathan, A. Srinivasan, A. Glynn, J. Eisenstein, and E. Gilbert, "You can't stay here: The efficacy of reddit's 2015 ban examined through hate speech," *ACM on Human-Computer Interaction*, vol. 1, no. CSCW, pp. 31:1–31:22, Dec. 2017.
- [3] J. Wilson, "Gab: alt-right's social media alternative attracts users banned from Twitter," <https://www.theguardian.com/media/2016/nov/17/gab-alt-right-social-media-twitter>, November 2016.
- [4] M. Wendling, "Gab: Free speech haven or alt-right safe space? By BBC Trending," <http://www.bbc.com/news/blogs-trending-38305402>, December 2016.
- [5] K. Shu, A. Sliva, S. Wang, J. Tang, and H. Liu, "Fake news detection on social media: A data mining perspective," *ACM SIGKDD Explorations Newsletter*, 2017.
- [6] W. Quattrociocchi, A. Scala, and C. R. Sunstein, "Echo chambers on facebook," 2016.
- [7] S. A. Munson, S. Y. Lee, and P. Resnick, "Encouraging reading of diverse political viewpoints with a browser widget." in *ICWSM*, 2013.
- [8] S. Park, S. Kang, S. Chung, and J. Song, "Newscube: delivering multiple aspects of news to mitigate media bias," in *SIGCHI*. ACM, 2009, pp. 443–452.
- [9] M. Babaei, J. Kulshrestha, A. Chakraborty, F. Benevenuto, K. P. Gummadi, and A. Weller, "Purple feed: Identifying high consensus news posts on social media," in *AIES*, 2018.
- [10] S. Zannettou, B. Bradlyn, E. De Cristofaro, M. Sirivianos, G. Stringhini, H. Kwak, and J. Blackburn, "What is gab? a bastion of free speech or an alt-right echo chamber?" *CyberSafety (WWW)*, 2018.
- [11] K. Popper and K. S. Popper, *The open society and its enemies*. Routledge, 2012.
- [12] D. Mitchell, "The liberalization of free speech: Or, how protest in public space is silenced," *Spaces of Contention: Spatialities and Social Movements*, vol. 47, 2016.
- [13] L. Silva, M. Mondal, D. Correa, F. Benevenuto, and I. Weber, "Analyzing the targets of hate in online social media," in *ICWSM*, 2016.
- [14] M. Mondal, L. A. A. Silva, and F. Benevenuto, "A measurement study of hate speech in social media," in *ACM HT*, 2017.
- [15] D. Bhattacharya and S. Ram, "Sharing news articles using 140 characters: A diffusion analysis on twitter," in *ASONAM*, 2012.
- [16] A. Chakraborty, S. Ghosh, N. Ganguly, and K. P. Gummadi, "Dissemination biases of social media channels: On the topical coverage of socially shared news," in *ICWSM*, 2016.
- [17] J. An, D. Quercia, M. Cha, K. Gummadi, and J. Crowcroft, "Sharing political news: the balancing act of intimacy and socialization in selective exposure," *EPJ Data Science*, 2014.
- [18] M. Hu, S. Liu, F. Wei, Y. Wu, J. Stasko, and K.-L. Ma, "Breaking news on twitter," in *CHI*. ACM, 2012.
- [19] T. Rodrigues, F. Benevenuto, M. Cha, K. Gummadi, and V. Almeida, "On word-of-mouth based discovery of the web," in *SIGCOMM*. ACM, 2011.
- [20] C. S. Lee and L. Ma, "News sharing in social media: The effect of gratifications and prior experience," *Computers in Human Behavior*, 2012.
- [21] K. Garimella, G. D. F. Morales, A. Gionis, and M. Mathioudakis, "Political discourse on social media: Echo chambers, gatekeepers, and the price of bipartisanship," in *WWW*, 2018.
- [22] J. Reis, H. Kwak, J. An, J. Messias, and F. Benevenuto, "Demographics of news sharing in the us twittersphere," in *HYPERTEXT*. ACM, 2017.
- [23] R. Ottoni, E. Cunha, G. Magno, P. Bernardina, W. Meira Jr, and V. Almeida, "Analyzing right-wing youtube channels: Hate, violence and discrimination," in *Proceedings of the 10th ACM Conference on Web Science*, 2018, pp. 323–332.
- [24] K. Flynn, "Tech is cracking down on hate speech—but it's still thriving on Gab and 4chan," <https://mashable.com/2017/08/17/alt-right-free-speech-online-network-gab/#EI0Lb40BQmqx>, August 2017.
- [25] J. Kulshrestha, M. Eslami, J. Messias, M. B. Zafar, S. Ghosh, K. P. Gummadi, and K. Karahalios, "Quantifying search bias: Investigating sources of bias for political searches in social media," *CSCW*, 2017.
- [26] T. Speicher, M. Ali, G. Venkatadri, F. N. Ribeiro, G. Arvanitakis, F. Benevenuto, K. P. Gummadi, P. Loiseau, and A. Mislove, "On the Potential for Discrimination in Online Targeted Advertising," in *FAT**, 2018.
- [27] J. Messias, P. Vikatos, and F. Benevenuto, "White, man, and highly followed: Gender and race inequalities in twitter," in *EEE/WIC/ACM*, 2017.
- [28] A. Chakraborty, J. Messias, F. Benevenuto, S. Ghosh, N. Ganguly, and K. P. Gummadi, "Who makes trends? understanding demographic biases in crowdsourced recommendations," in *ICWSM*, 2017.
- [29] Southern Poverty Law Center, "Extremists," <https://www.splcenter.org/fighting-hate/extremist-files/individual>, 2018, (Accessed on 01/25/2018).
- [30] Anti-Defamation League, "From alt right to alt lite: Naming the hate," <https://www.adl.org/education/resources/backgrounders/from-alt-right-to-alt-lite-naming-the-hate>, 2017.
- [31] C. Budak, S. Goel, and J. M. Rao, "Fair and balanced? quantifying media bias through crowdsourced content analysis," *Public Opinion Quarterly*, 2016.
- [32] A. Mitchell, J. Gottfried, J. Kiley, and K. E. Matsa, "Political polarization & media habits," *Pew Research Center*, 2014.
- [33] E. Bakshy, S. Messing, and L. A. Adamic, "Exposure to ideologically diverse news and opinion on facebook," *Science*, vol. 348, no. 6239, pp. 1130–1132, 2015.
- [34] L. Peer and T. B. Ksiazek, "Youtube and the challenge to journalism: new standards for news videos online," *Journalism Studies*, vol. 12, no. 1, pp. 45–63, 2011.
- [35] G. Hanson and P. Haridakis, "Youtube users watching and sharing the news: A uses and gratifications approach," *Journal of Electronic Publishing*, vol. 11, no. 3, 2008.
- [36] D. Antoniadis, I. Polakis, G. Kontaxis, E. Athanasopoulos, S. Ioannidis, E. P. Markatos, and T. Karagiannis, "we. b: The web of short urls," in *WWW*, 2011.
- [37] J. Reis, F. Benevenuto, P. V. de Melo, R. Prates, H. Kwak, and J. An, "Breaking the news: First impressions matter on online news," in *ICWSM*, 2015.