

An Analysis of Interactions Within and Between Extreme Right Communities in Social Media

Derek O’Callaghan¹, Derek Greene¹, Maura Conway², Joe Carthy¹, Pádraig Cunningham¹

¹ School of Computer Science & Informatics, University College Dublin,
{derek.ocallaghan,derek.greene,joe.carthy,padraig.cunningham}@ucd.ie

² School of Law & Government, Dublin City University,
maura.conway@dcu.ie

Abstract. Many extreme right groups have had an online presence for some time through the use of dedicated websites. This has been accompanied by increased activity in social media websites in recent years, which may enable the dissemination of extreme right content to a wider audience. In this paper, we present exploratory analysis of the activity of a selection of such groups on Twitter, using network representations based on reciprocal follower and mentions interactions. We find that stable communities of related users are present within individual country networks, where these communities are usually associated with variants of extreme right ideology. Furthermore, we also identify the presence of international relationships between certain groups across geopolitical boundaries.

Keywords: network analysis, social media, community detection, Twitter, extreme right

1 Introduction

Groups associated with the extreme right have maintained an online presence for some time [1], where dedicated websites have been employed for the purposes of content dissemination and member recruitment. Recent years have seen increased activity by these groups in social media websites, given the potential to access a far wider audience than was previously possible. In this paper, we present exploratory analysis of the activity of a selection of these groups on Twitter, where the focus is upon groups of a fascist, racist, supremacist, extreme nationalist or neo-Nazi nature, or some combination of these. Twitter’s features enable extreme right groups to disseminate hate content with relative ease, while also facilitating the formation of communities of users around variants of extreme right ideology. Message posts (*tweets*) by such groups are often used to redirect users to content hosted on external websites, for example, dedicated websites managed by particular groups, or content sharing websites such as YouTube.

The first objective of this analysis is the detection of communities of users associated with extreme right groups within individual countries, using network

representations of user interactions. For the purpose of this exploratory work, we have retrieved Twitter data for a selection of eight countries. Having ranked communities detected within each country network representation based on their stability, we generate corresponding descriptions that may be used to provide an interpretation of the underlying community ideology. Our second objective is the identification of international relationships between certain groups that transcend geopolitical boundaries, using two network representations of the interactions between the core users from the eight country sets. It appears that a certain number of international relationships exist, where linguistic and geographical proximity are highly influential.

In Section 2, we provide a description of related work based on the online activities of extremist groups. The generation of the Twitter data sets is then discussed in Section 3. Next, in Section 4, we describe the detection of local extreme right communities within individual countries using two case studies based on the USA and Germany networks. Analysis of the international relationships between extreme right groups from the eight countries is presented in Section 5. Finally, conclusions and suggestions for future work are discussed in Section 6.

2 Related Work

The online activities of different varieties of extremist groups including those associated with the extreme right have been the subject of a number of studies. For example, Burris et al. [1] used social network analysis to study a network based on the links between a selection of white supremacist websites. Chau and Xu [2] studied networks built from users contributing to hate group and racist blogs. The potential for online radicalization through exposure to jihadi video content on YouTube was investigated by Conway and McInerney [3].

Tateo [4] analyzed groups associated with the Italian extreme right using networks based on links between group websites. Caiani and Wagemann [5] studied similar Italian groups along with those from the German extreme right. In their review of the conservative movement in the USA, Blee and Creasap [6] discuss the engagement in online activity as part of an overall mobilization strategy by the more extremist groups within it. Simi [7] analyzed the apparent decentralization of white supremacist groups according to, among other aspects, their online activity. As the majority of this related work involved the study of dedicated websites managed by extreme right groups, we felt that an analysis of their activity in social media would complement this by providing additional insight into the overall online presence of these groups.

3 Data Retrieval with User Curation

Twitter data was collected to facilitate the analysis of contemporary activity by extreme right groups. As the hypothesis was that extreme right communities within social media would tend to be relatively smaller than mainstream communities, the data was sampled using *core sets* of relevant user accounts, one set

per country of interest. A description of this process is available in the extended version of this paper [8].

4 Local Community Detection

The interactions between users within the individual country sets were analyzed with the objective of detecting communities of related users. At the country level, an interaction is defined as one user “mentioning” another with the inclusion of “@<username>” within a tweet, where reciprocal mentions between users can be considered as a dialogue [9], thus potentially indicating the presence of a stronger relationship. A network is created with n nodes representing users, and undirected weighted edges representing reciprocal mentions between pairs of users, with weights based on the number of mentions. Currently, all mentions occurrences are considered rather than selecting those from a specific time period. Any connected components of size < 10 are filtered at this point.

We use the method of Greene et al. [10], which is a variant of the recent work by Lancichinetti & Fortunato [11], to generate a set of stable *consensus communities* from a reciprocal mentions network, where 100 runs of the OSLOM algorithm [12] are used to generate the consensus communities. In this analysis, we are primarily concerned with the detection of the most “significant” communities with the strongest signals in the network. Therefore, the consensus communities are ranked based on the stability of their members with respect to the corresponding consensus matrix. We employ the widely-used adjustment technique introduced by [13] to correct for chance agreement:

$$\text{CorrectedStability}(C) = \frac{\text{Stability}(C) - \text{ExpectedStability}(C)}{1 - \text{ExpectedStability}(C)} \quad (1)$$

A value close to 1 will indicate that C is a highly-stable community. As higher values of the threshold parameter τ used with the consensus method result in sparser consensus networks, and having tested with values of τ in the range [0.1, 0.8], we selected $\tau = 0.5$ as a compromise between node retention and more stable communities. Finally, descriptions are generated using a TF-IDF vector for each community, where the terms are hashtags contained within tweets posted by users within the community. A description consists of the top ten hashtag terms ranked using their TF-IDF values. Details of the consensus communities found in the USA and Germany case study networks can be found in Table 1.

Table 1. Consensus communities in local reciprocal mentions networks.

<i>Country</i>	<i>Original network</i>		<i>Consensus network</i>		<i>Communities</i>	<i>Core users</i>
	<i>Nodes</i>	<i>Edges</i>	<i>Nodes</i>	<i>Edges</i>		
USA	835	2501	672	6876	55	29
Germany	247	646	167	799	18	46

4.1 Case Study: USA

Table 2. USA reciprocal mentions - selected consensus communities ($\tau = 0.5$)

<i>Communities with core users</i>				
Id	Description	Size	Core (%)	Score
A	aryan, pipa, thewhiterace, shtf, wpww, masterrace, londonriots, ukriots, wwii, nazi	11	2 (18%)	0.89
B	wpww, tcot, gop, truth, teaparty, trayvontruth, blackpower, sopa, mlk, treyvon	24	2 (8%)	0.81
C	ubnp, wpradio, contest, wpww, pandora, wpwwgiveaway, nevershouldyouever, hch, rahowa, survival	75	6 (8%)	0.81
D	prayforthis, glory, genocide, diversity, africans, antiracist, equality, asia, asians, mustsee	11	3 (27%)	0.81
<i>Communities without core users</i>				
Id	Description	Size	Core (%)	Score
E	zumaspear, da, sa, afrikaans, anc, ancyl, zuma, southafrica, malema, afrikaners	38	0 (0%)	0.91
F	rochdale, brighton, edl, mfe, uaf, islam, luton, dewsbury, bbcsm, labour	33	0 (0%)	0.85

A selection of relevant communities having high stability scores can be found in Table 2. Community A would appear to be national socialist/white power in nature, with the appearance of hashtags such as “aryan”, “thewhiterace”, “wpww” (white pride world wide), “masterrace” and “nazi”. An analysis of the users and associated profiles finds references to the *American Nazi Party*, along with other related terms such as *14* (a reference to a 14-word slogan coined by the white supremacist David Lane), and *88* (“H” is the 8th letter in the alphabet, and *88* signifies “heil hitler”) in user names. Users appear to be mostly from the USA, although a small number of European users are also present (“londonriots”, “ukriots”). Some similar white power themes (e.g. “rahowa” - RAcial HOLy WAR) appear in communities B and C, where both communities contain a number of North American users including some who are promoting external white power radio station websites.

Most of the users in community D (both North American and European) appear to be connected with a number of *white rabbit* websites which allege the existence of “white genocide”. Communities E and F are interesting as neither community contains a single user from the USA core set. However, an analysis of the users in both communities shows that reciprocal follower relationships with USA users are common. The former appears to consist of white South African users where some profiles contain racist and national socialist references, while the latter contains many references to the *English Defence League* (EDL), a group opposed to the alleged spread of militant islamism within the UK.

4.2 Case Study: Germany

A selection of relevant communities having high stability scores can be found in Table 3. An analysis of the users in community A finds them to be primarily associated with the town of Geithain, near Leipzig in Sachsen (“geithainer”, “geithain”, “gha” and “lvz” hashtags). Users belonging to various extreme right groupings are present, such as *Freies Netz* (FN - “information portal” websites hosting related content) and the *Junge Nationaldemokraten* (JN -

Table 3. Germany reciprocal mentions - selected consensus communities ($\tau = 0.5$)

<i>Communities with core users</i>				
Id	Description	Size	Core (%)	Score
A	bollywood, lvz, geithainer, geithain, tdi, unsterblichen, gha, mephisto, imc, jcsyhra	10	3 (30%)	0.89
B	saalfeld, gera, pcrecords, apw, 13februar, volkstod, otz, geraer, altenburger, rfd	14	5 (36%)	0.79
C	130abschaffen, unbrennt, raz09, mobilisierungsvideo, golf, demokraten, spreelichter, meilederdemokratie, sfb, halloween	10	2 (20%)	0.77
D	guben, spreelichter, apw, jingle, weltanschauung, vetschau, bock, 17august, podcast, altermedia	24	10 (42%)	0.72
<i>Communities without core users</i>				
Id	Description	Size	Core (%)	Score
E	bamberg, flugblatt, ovg, trke, stolberg, ruhrgebiet, wattenscheid, vorstellungsflugblatt, rat, landesgartenschau	8	0 (0%)	0.94

Young National Democrats, youth wing of the extreme right *Nationaldemokratische Partei Deutschlands - National Democratic Party of Germany - NPD*). The “unsterblichen” (immortals) hashtag refers to anti-democratic flashmob marches that were occurring sporadically throughout Germany. These protests were linked to *Spreelichter*, an extreme right group that was recently banned by the local authorities. who used social media to propagate national socialist-related material.

Community B also appears to be mostly related to a geographical location, namely the federal state of Thüringen (“altenburg”, “gera”, “saalfeld”). The “rfd” hashtag refers to “Rock für Deutschland”; a concert organised by the NPD in Gera. Other relevant hashtags include “apw” (außerparlamentarischer Widerstand - non-parliamentary resistance), “13februar” (1945 bombing of Dresden by Allied forces) and “volkstod” (perceived destruction of German people and traditions). The theme of community C is related to the “130abschaffen” hashtag, which refers to demands for the abolition of a paragraph in the German penal code associated with the criminalization of incitement to hatred, along with denial and/or justification of the Holocaust and national socialist rule.

Community D contains a wide range of users from groups such as FN, JN and *Spreelichter/Unsterblichen*. These users would appear to be quite active, with many tweets containing URLs linking to content hosted on external websites such as YouTube or other dedicated websites. No core users are present in community E, but all members appear to be related to the NPD. The hashtags “bamberg” and “ruhrgebiet” refer to locations within Germany, and the presence of other hashtags such as “flugblatt” (flyer/leaflet/pamphlet) and “rat” (council/councillor) along with separate analysis of the tweet content may indicate mobilization prior to elections.

Some further details of the communities from both case studies can be found in the extended version of this paper [8].

5 International Relationships

We also analyzed the international relationships between the various groups within the data sets, based on the interactions between the core users from the

eight country sets. Two types of undirected network were generated; a followers network consisting of user nodes and unweighted edges representing follower links between users, and a mentions network similar to those described in the previous section. As with the country-based networks, only reciprocal edges were used in order to capture stronger relationships, all stored follower and mentions instances were included, and connected components of size < 10 were filtered.

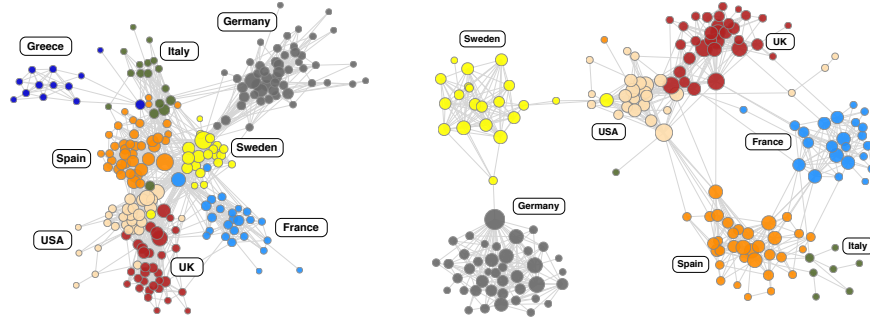


Fig. 1. International core networks: followers (left - 226 nodes, 1603 edges) and mentions (right - 184 nodes, 856 edges). Node size is proportional to degree.

5.1 International Follower Awareness

The international followers network can be seen in Fig. 1 (left). As might be expected, most of the follower relationships are between users from the same country, although a certain number of international relationships are identifiable. It would appear that linguistic and geographical proximity is influential here; similar behaviour with respect to social ties in Twitter was identified by Takhteyev et al. [14]. However, there appear to be some exceptions to the influence of geographical proximity, most notably, Swedish (yellow) and Italian (green) users that are not co-located with their respective country nodes. In both cases, the majority of tweets from these users are in English, which presumably ensures a wider audience.

From an analysis of other central nodes in the network (using betweenness centrality), it would seem that those involved in the dissemination of material (text, audio, video) with the use of external URLs, or media representatives such as extreme right news websites and radio stations, are attempting to raise awareness amongst a variety of international followers. This is especially the case when the English language is used.

5.2 International Dialogue

It is reasonable to consider that the follower-based networks represent passive relationships when compared with the mentions-based networks, where links can

represent actual dialogue between users. The mentions network in Fig. 1 (right) can be seen to be somewhat smaller than the corresponding followers network in Fig. 1 (left); for example, none of the Greek core users are present. Apart from this, the network has a similar structure to that of the followers network, in that most interaction occurs within individual country-based communities. Connections between these communities do exist, but are fewer than in the followers network. The influence of linguistic proximity appears to take precedence here, with the use of English playing a major role as mentioned in the previous section. In the case of the German community, while the followers network contains a number of connections with other international users, this has now been reduced to a single connection with a user acting as an English-language Twitter channel for a Swedish nationalist group. Similarly, the Swedish user co-located with the USA community is the same user as that in the followers network, who appears to be involved in many English-based dialogues.

It should also be emphasized at this point that this analysis does not necessarily provide extensive coverage of the international relationships between all extreme right groups that are active on Twitter. As it is possible that the data retrieved from Twitter is incomplete, the objective is to demonstrate the existence of these relationships by means of an exploratory analysis. This caveat also applies to the country-based community analysis.

6 Conclusions and Future Work

Extreme right groups have become increasingly active in social media websites such as Twitter in recent years. We have presented an exploratory analysis of the activity of a selection of such groups using network representations based on reciprocal follower and mentions interactions. The existence of stable communities of associated users within individual countries has been demonstrated, and we have also identified international relationships between certain groups across geopolitical boundaries. Although a certain awareness exists between users based on the follower relationship, it would appear that mentions interactions indicate stronger relationships where linguistic and geographical proximity are highly influential, in particular, the use of the English language. In relation to this, media user accounts such as those associated with extreme right news websites and radio stations, along with external websites hosting content such as music or video, are a popular mechanism for the dissemination of material among users from a variety of disparate groups.

Although some of the detected communities can be associated with a specific extreme right group or ideology, this is less clear in other cases where communities appear to contain members from a variety of known groups. This may be a consequence of incompleteness in the data sets retrieved for this analysis. It may also be related to variances in Twitter usage patterns in different countries. The laws of different countries should also be taken into consideration, as an opinion that may be legally voiced in one country may not be permitted in another, particularly within the context of extreme right ideals. However, it may also be the

case that social media websites are merely used by such groups to disseminate related material to a wider audience, with the majority of subsequent interaction occurring elsewhere.

In future work, we will address the issue of incompleteness in the data sets, including the current disparity in core set sizes. Local community analysis of countries other than the USA and Germany will be performed. We also plan to study the temporal properties of these networks which will provide insight into the evolution of extreme right communities over time.

Acknowledgements. This research was supported by 2CENTRE, the EU funded Cybercrime Centres of Excellence Network and Science Foundation Ireland Grant 08/SRC/I1407 (Clique: Graph and Network Analysis Cluster).

References

1. Burris, V., Smith, E., Strahm, A.: White Supremacist Networks on the Internet. *Sociological Focus* **33**(2) (2000) 215–235
2. Chau, M., Xu, J.: Mining communities and their relationships in blogs: A study of online hate groups. *Int. J. Hum.-Comput. Stud.* **65**(1) (January 2007) 57–70
3. Conway, M., McInerney, L.: Jihadi Video and Auto-radicalisation: Evidence from an Exploratory YouTube Study. In: *Proceedings of the 1st European Conference on Intelligence and Security Informatics. EuroISI '08, Berlin, Heidelberg, Springer-Verlag* (2008) 108–118
4. Tateo, L.: The Italian Extreme Right On-line Network: An Exploratory Study Using an Integrated Social Network Analysis and Content Analysis Approach. *Journal of Computer-Mediated Communication* **10**(2) (2005) 00–00
5. Caiani, M., Wagemann, C.: Online Networks of the Italian and German Extreme Right. *Information, Communication & Society* **12**(1) (2009) 66–109
6. Blee, K.M., Creasap, K.A.: Conservative and Right-Wing Movements. *Annual Review of Sociology* **36** (2010) 269–286
7. Simi, P.: Why Study White Supremacist Terror? A Research Note. *Deviant Behaviour* **31**(3) (2010) 251–273
8. O’Callaghan, D., Greene, D., Conway, M., Carthy, J., Cunningham, P.: An Analysis of Interactions Within and Between Extreme Right Communities in Social Media (extended version). *ArXiv e-prints* (<http://arxiv.org/abs/1206.7050>) (June 2012)
9. Macskassy, S.: On the Study of Social Interactions in Twitter. In: *Sixth International AAAI Conference on Weblogs and Social Media. ICWSM* (2012)
10. Greene, D., O’Callaghan, D., Cunningham, P.: Identifying Topical Twitter Communities via User List Aggregation. *ArXiv e-prints* (<http://arxiv.org/abs/1207.0017>) (June 2012)
11. Lancichinetti, A., Fortunato, S.: Consensus clustering in complex networks. *Sci. Rep.* **2** (03 2012)
12. Lancichinetti, A., Radicchi, F., Ramasco, J., Fortunato, S., Ben-Jacob, E.: Finding statistically significant communities in networks. *PLoS ONE* **6**(4) (2011) e18961
13. Hubert, L., Arabie, P.: Comparing partitions. *Journal of Classification* (1985) 193–218
14. Takhteyev, Y., Gruzd, A., Wellman, B.: Geography of Twitter networks. *Social Networks* **34**(1) (2012) 73 – 81