

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

Derek O’Callaghan<sup>1</sup>, Derek Greene<sup>1</sup>, Maura Conway<sup>2</sup>, Joe Carthy<sup>1</sup>, Pádraig Cunningham<sup>1</sup>

**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 [1,2], 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, to which access is usually unrestricted, 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. These posts also assist in the mobilization of participants for subsequent interactions such as offline demonstrations or other offline activities.

Our investigation has found that communities of users associated with extreme right groups are indeed present on Twitter. The primary objective of this analysis is the detection of such communities within individual countries, using network representations of user interactions. In this context, an interaction is defined as one user “mentioning” another within a tweet, where reciprocal mentions between users can be considered as a dialogue [3], potentially indicating the presence of a stronger relationship. For the purpose of this exploratory work, we have retrieved Twitter data, including profile information and posted tweets, for a selection of eight countries, using a sampling method that requires the identification of sets of *core* users who are considered

to be highly relevant. A network representation is extracted for each of these country data sets, from which communities of related users are detected. Having ranked these communities based on their stability, we can describe them using the *hashtags* contained within tweets posted by the member users. Each such description can then be used in conjunction with manual analysis of the user profiles, tweets and external websites to provide an interpretation of the underlying community ideology.

Our secondary objective is the identification of international relationships between certain groups that transcend geopolitical boundaries. This involves the analysis of two network representations of the interactions between the core users from the eight country sets, using the follower relationships and mentions dialogues. It appears that a certain amount of international awareness exists between users based on the follower relationship, while mentions interactions indicate stronger relationships 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, using profile data and tweets for users from each country of interest, is then discussed in Section 3. Next, in Section 4, we describe the detection of local extreme right communities within individual countries, including the methodology used for network derivation, community detection, stability ranking and description generation. In this exploratory analysis, we focus on two case studies, using the USA and Germany networks, where we offer an interpretation of a selection of these communities. Due to the sensitivity of the subject matter, and in the interests of privacy, individual users are not identified; instead, we restrict discussion to known extreme right groups and their affiliates. Analysis of the international relationships between extreme right groups from the eight countries is presented in Section 5. Finally, the overall conclusions are discussed in Section 6, and some suggestions for future work are made.

## 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. Burris et al. [2] used social network analysis to study a network based on the links between a selection of white supremacist websites. They found this network to be relatively decentralized with multiple centres of influence, while also appearing to be mostly undivided along doctrinal lines. Similar decentralization and multiple communities were found by Chau and Xu [4] in their study of networks built from users contributing to hate group and racist blogs. They also found that some of these groups exhibited transnational characteristics. The potential for online radicalization through exposure to jihadi video content on YouTube was investigated by Conway and McInerney [5]. A continuation of this work by Bermingham et al. [6] included network analysis of the associated YouTube users, where it was suggested that a potentially increased online leadership role may be attributed to users claiming to be women, according to centrality, network density and average speed of communication. Sureka et al. [7] also studied the activity of extremist users within YouTube, investigating content properties along with hidden network communities.

In a similar approach to that of Burris et al. , Tateo [8] analyzed groups associated with the Italian extreme right, using networks based on links between group websites. A combination of social network analysis and content analysis methods were used. Caiani and Wagemann [9] studied similar Italian groups along with those from the German extreme right. They found the German network to be structurally centralized to a greater extent than that of the Italian groups. A related study by Caiani and Parenti [10] concluded that Spanish extreme right groups do not exploit the Internet for the purposes of communication and mobilization to the same degree as other countries. The contents of websites belonging to central nodes within Russian extreme right networks were analyzed by Zuev [11]. In their review of the conservative movement in the USA, Blee and Creasap [12] discuss the engagement in online activity as part of an overall mobilization strategy by the more extremist groups within it. McNamee et al. [13] analyzed the messages from a number of hate group websites, characterizing them using four themes; education, participation, invocation and indictment. Simi [14] suggested that the apparent decentralization of white supremacist groups according to, among other aspects, their online activity, should not be confused with disorganization or irrelevance. 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, a form of snowball sampling was applied rather than using a random sampling approach. One of the objectives of this analysis was the detection of extreme right communities within individual countries. Given this, we used the *curation* method as proposed by Greene et al. [15] to generate *core sets* of relevant user accounts, one set per country of interest. This involved the initial identification of country seed user sets, where the user accounts within these seed sets were chosen based on a number of criteria (see below). Prior knowledge of extreme right groups informed the selection of countries of interest.

Following the *bootstrapping phase*, the core sets were expanded over a number of iterations based on manual analysis of the users proposed in the *recommendation phases* (details of the final core sets can be found in Table 1). A selection of recommended users were added to the core sets using the following criteria:

- ⤴ Relevant profiles; for example, those containing references to known groups or employing extreme right symbols.
- ⤴ Recent tweet activity (within the last three months).
- ⤴ Similar YouTube user account profiles.
- ⤴ Follower relationships with known relevant users.

- ⤴ Users with self-curated lists (Twitter feature) containing relevant users.
- ⤴ Extreme right media accounts; for example, record labels, online music stores, radio stations, concert organisers etc.

**Table 1:** Sizes of core user sets for eight countries of interest.

<i>Country</i>	<i>Number of Core Users</i>
France	25
Germany	53
Greece	14
Italy	17
Spain	42
Sweden	21
UK	32
USA	32

A number of recommended users were also ignored, such as inactive users, or those that were not deemed to be related to the extreme right. These included traditional conservative (e.g. centre-right) users, non-conformists/anti-establishment users considered to be left-wing, and conspiracy theorists. As the focus of the analysis was on extreme right groups, higher-profile politicians or political parties were ignored for the most part, with a minor number of these users included where it was felt that there was a close association with relevant users. Although expansion of the core sets using this criteria required a certain amount of subjectivity, it is felt that this was acceptable given that the requirement was to identify a relatively small number of relevant user accounts, and relevance is often immediately evident from manual profile analysis. A major obstacle was the language barrier, as the use of online translation tools were not always helpful in the case of ambiguous profiles. In such cases where the relevance of a user profile was inconclusive, the user was ignored. Although this leads to a disparity in core set size, for example, Spain and Germany have a relatively larger number of users, core set size does not necessarily influence the size of the communities detected within the generated networks, as demonstrated later in the case study sections.

Profile data including followers, friends, tweets and list memberships were retrieved for each of the core users, as limited by the current Twitter API restrictions. The bootstrap process also retrieves a selection of profile data of those users having a follower relationship with a core user. In order to address possible incompleteness, additional profile data was also retrieved for any users having a reciprocal follower relationship with more than one core user, to produce the final data sets for analysis. As the Twitter follower relationship tends to exhibit lower reciprocity than other social networking sites [16], the understanding was that this action would be largely isolated to those users having a relatively stronger relationship with core users.

## 4 Local Community Detection

Having retrieved the user data from Twitter, 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 [3], 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 occurrences. 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.

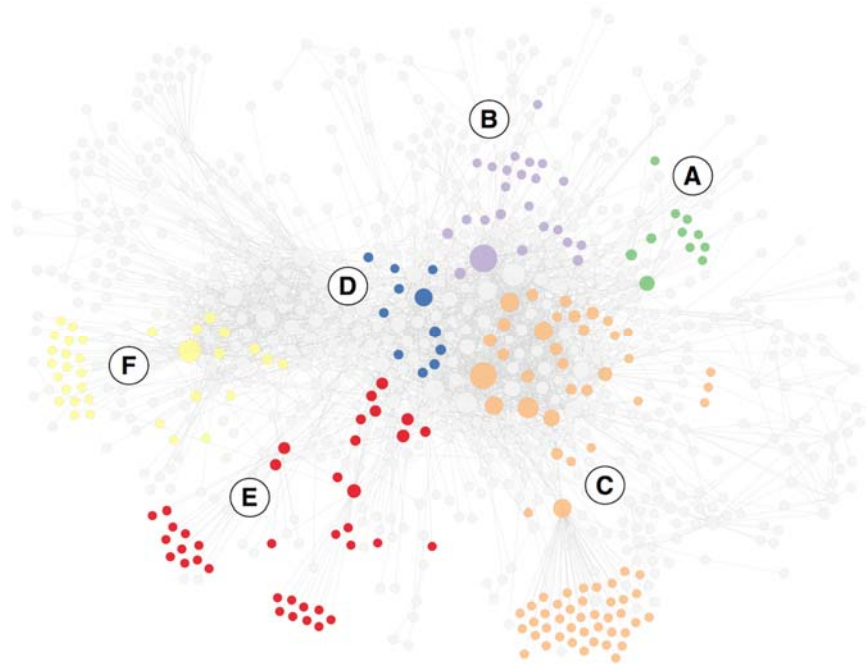
We use the method of Greene et al. [17], which is a variant of the recent work by Lancichinetti & Fortunato [18], to generate a set of stable *consensus communities* from a reciprocal mentions network, where 100 runs of the OSLOM algorithm [19] are used to generate the consensus communities. In this 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 [20] 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  $\tau$  used with the consensus method result in sparser consensus networks, and having tested with values of  $\tau$  in the range  $[0.1, 0.8]$ , we selected  $\tau = 0.5$  as a compromise between node retention and more stable communities. Finally, community 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.

### 4.1 Case Study: USA

A reciprocal mentions network (Fig. 1) was created for the USA core set and associated data, consisting of 835 nodes and 2501 edges, where 29 of the original 32 core set nodes were present. The consensus communities were then generated using the methodology described earlier. The consensus network generated with threshold  $\tau = 0.5$  consisted of 672 nodes and 6876 edges. Fifty-five communities were detected, and a selection of relevant communities having high stability scores can be found in Table 2.



**Fig. 1:** USA reciprocal mentions network - 835 nodes and 2501 edges, selected consensus communities are labelled. Node size is based on degree.

**Table 2:** USA reciprocal mentions - selected consensus communities (consensus network threshold  $\tau = 0.5$ )

<i>Communities with core users</i>					
Id	Colour	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	Colour	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,	33	0 (0%)	0.85

		dewsbury, bbcsm1, labour			
--	--	--------------------------	--	--	--

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. There are also references to skinhead groups, including a website where related media and merchandise can be found for sale. Users appear to be mostly from the USA, although a small number of European users are also present. References to the London riots of 2011 are also made within tweets, where blame has been apportioned to non-whites. “Pipa” refers to the “PROTECT IP Act”, which appears to be a shared concern among such groups with respect to potential curbs to civil liberties.


Some similar white power themes appear in community B, whose description also includes “wpww”, although they would appear to be more subtle than those of community A. An analysis of the user profiles finds a certain number of North American white power users with some similar references as community A, albeit on a smaller scale. One of the core users promoting a radio station website could be considered influential here. The “trayvontruth” and “treylon” hashtags refer to the recent fatal shooting in Florida of Trayvon Martin, a 17-year-old African-American teenager. It appears that this story is being used to propagate a message of alleged persecution against whites. References to “tcot” (top conservatives on Twitter) and “gop” (Grand Old Party aka the USA Republican Party) can also be seen, indicating the presence of more traditional conservative users. However, this does not necessarily point to any official link between these groups. The appearance of the “sopa” (Stop Online Piracy Act) hashtag is most likely analogous to that of “pipa” in community A’s description.

The users in community C form quite a large North American community which is primarily white power oriented, as can be seen from the frequent occurrences of “wpww” among the description hashtags, along with “rahowa” (Racial HOLy WAR). There are many national socialist references within the user profiles, and as with community A, many user names containing *14* and *88*. In a similar fashion to community B, an external white power radio station website seems to be influential here, which correlates with some of the hashtags such as “wpradio”, “contest” and “wpwwgiveaway”. 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”. Here, “antiracist” refers to the associated slogan “anti-racist is just a code word for anti-white”. This community also contains one liberal user with a high-level of posting activity who has interacted with a certain number of the other white power users.

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 users in community E appear to be white South Africans, with some profiles containing racist and national socialist references. Many tweets from these users relate to perceived cultural threats from black South Africans that

are often retweeted by international users. Almost all of the description hashtags are related to South Africa, such as “anc” (*African National Congress*, the current governing party) and “afrikaners” (primarily white ethnic group descended from European settlers). Community F consists mostly of users associated with the *English Defence League* (“edl” hashtag), a group opposed to the alleged spread of militant islamism within the UK. Hashtags such as “rochdale”, “brighton”, “luton” and “dewsbury” are UK locations, most likely indicating various EDL subgroups, while “uaf” refers to the *Unite Against Fascism* group; a staunch opponent of the EDL. Although both communities E and F would appear to have a certain amount of awareness and interaction with other nationalities as indicated by the follower relationships, the majority of interaction seems to be at the local level if mentions interactions are assumed to represent stronger relationships.

## 4.2 Case Study: Germany



**Fig. 2:** Germany reciprocal mentions network - 247 nodes, 646 edges, selected consensus communities are labelled. Node size is based on degree.

A reciprocal mentions network (Fig. 2) was created for the Germany core set and associated data, consisting of 247 nodes and 646 edges, where 46 of the original 53 core set nodes were present. The consensus communities were then generated using the methodology described earlier. The consensus network generated with threshold  $\tau=0.5$  consisted of 167 nodes and 799 edges. Eighteen communities were detected, and a selection of relevant communities having high stability scores can be found in Table 3.



**Table 3:** Germany reciprocal mentions - selected consensus communities (consensus network threshold  $\tau = 0.5$ )

<i>Communities with core users</i>					
Id	Colour	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	Colour	Description	Size	Core (%)	Score
E	■	bamberg, flugblatt, ovg, trke, stolberg, ruhrgebiet, wattenscheid, vorstellungsflugblatt, rat, landesgartenschau	8	0 (0%)	0.94

An analysis of the users in community A finds them to be primarily associated with the town of Geithain, near Leipzig in Sachsen. This can also be seen from hashtags such as “geithainer”, “geithain” and “gha”, along with the Leipzig-related hashtag “lvz” (Leipziger Volkszeitung - newspaper). Users belonging to various extreme right groupings are present, such as *Freies Netz* (“information portal” websites hosting related content) and the *Junge Nationaldemokraten* (Young National Democrats, youth wing of the extreme right *Nationaldemokratische Partei Deutschlands - National Democratic Party of Germany - NPD*). References to *Aktionsbüros* (coordination of activist activities) are also made. The “bollywood” hashtag refers to a Geithain business that has been subjected to repeated attacks by neo-Nazis, while “unsterblichen” (immortals) refers to anti-democratic flashmob marches that have been occurring sporadically throughout Germany. These protests are linked to *Spreelichter*, an extreme right group from Südbrandenburg that was recently banned by the local authorities.<sup>1</sup> They used social media to propagate national socialist-related material, including professional-quality videos of the marches themselves.

<sup>1</sup> [http://blog.zeit.de/stoerungsmelder/2012/06/19/das-ende-der-nazi-masken-show\\_8923](http://blog.zeit.de/stoerungsmelder/2012/06/19/das-ende-der-nazi-masken-show_8923)

Like community A, the users in community B appear to be mostly related to a geographical location, namely the federal state (Bundesland) of Thüringen, where the towns of Altenburg, Gera and Saalfeld can be found. Most of the users are associated with *Freies Netz*, with a number from other parts of Germany. The “rfd” hashtag refers to “Rock für Deutschland” - a concert organised by the NPD in Thüringen which will take place in July 2012. Other relevant hashtags include “apw” (außerparlamentarischer Widerstand - non-parliamentary resistance) and “13februar”, a reference to the bombing of Dresden which began on February 13, 1945. This event is usually commemorated by extreme right groups each year. Also relevant is “volkstod” (death of the people/nation), which refers to the perceived destruction of German people and traditions since World War II, often mentioned in *Spreelichter*-related content.


The binding 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. Although other relevant hashtags such as “apw”, “volkstod” and “spreelichter” are present, some of the users have been inactive for some time, and it appears that these may have been replaced with new accounts. Other interesting hashtags are “mobilisierungsvideo” (mobilization video) and “unibrennt”, where the latter refers to a student protest movement. This movement appears to be unrelated, but the use of this hashtag by known extreme right users may be strategic as it could ensure the propagation of tweets to a wider audience.

Community D is the largest of the selected consensus communities, and contains a wide range of users from groups such as *Aktionsbüros/bündnis*, *Freies Netz*, *Junge Nationaldemokraten* and *Spreelichter/Unsterblichen*. These users would appear to be quite active, with many “apw”-related tweets, along with tweets containing URLs linking to content hosted on external websites such as YouTube or other dedicated websites. The “altermedia” hashtag refers to a collective of politically-incorrect/nationalist-oriented news websites, where the German website contains content such as *Spreelichter* articles and a section related to the NPD. “Guben” is a town in the state of Brandenburg on the border between Germany and Poland, and its many occurrences may indicate a geographical connection with Brandenburg, particularly as *Spreelichter* was also based in this state. It also appears that some of the users here are replacement accounts for older versions that are members of community C.

Finally, as with the USA case study, a community without any core set users has been selected for analysis. In this case, all of the users in community E 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. Separately, the “türke” hashtag refers to people of Turkish descent, and is often used in tweets from community E (and other) users alleging the involvement of such people in criminal activity; a common accusation by German extreme right groups. Similarly, “stolberg” refers to the town where a German teenager was killed by non-Germans in 2008, which is also the focus of an annual extreme right commemoration.

## 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 (Table 1). 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 where edges were created for each instance of a user “mentioning” another with the inclusion of “@<username>” within a tweet. 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. 3:** International core followers network (reciprocal) - 226 nodes, 1603 edges. Node size is based on degree.

### 5.1 International Follower Awareness

The international followers network can be seen in Fig. 3. 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

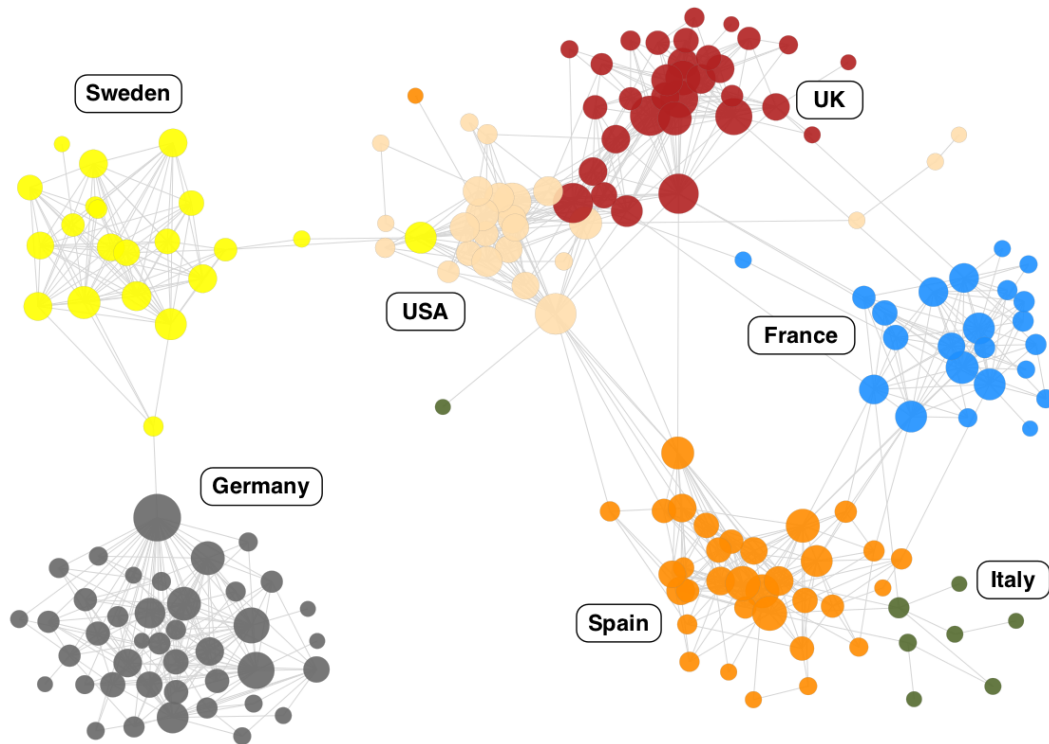
proximity is influential here, for example, there appears to be a close relationship between the Spanish and Italian (and to a lesser extent, French) users, with strong connections also between the UK and USA. Similar behaviour with respect to social ties in Twitter has been identified by Takhteyev et al. [21] and Kulshrestha et al. [22]. 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. The former user is a Swedish representative of a pan-Scandinavian group espousing national socialist ideals, who appears to be interacting with many international users, particularly from the USA. This user has been increasingly active in recent months.

The Italian user is a national socialist whose tweets often contain URLs to music or video content hosted on external websites, but it is unclear if a direct association exists with any particular group. 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

The follower-based relationship between international users could be considered as passive when compared with that of the mentions-based networks, where such interactions can represent actual dialogue between users. The mentions network in Fig. 4 can be seen to be somewhat smaller than the corresponding followers network in Fig. 3; 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. For example, a relatively large number of connections remain present between the UK and USA users. 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 with international users.

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.



**Fig. 4:** International core mentions network (reciprocal) - 184 nodes, 856 edges. Node size is based on degree.

## 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 ideals among users from a variety of disparate groups.

Although a certain number of the detected communities can be associated with a specific extreme right group or ideology, this is more ambiguous 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, for example, the use of Twitter tends to be more

prevalent in the UK and the USA than in other countries such as Germany<sup>2</sup>. 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, and we will collect data beyond the eight countries used in this analysis. We also plan to study the temporal properties of these networks which will provide insight into the evolution of extreme right communities over time.

## References

1. Hoffman, D.S.: The Web of Hate: Extremists Exploit the Internet. Anti-Defamation League (1996)
2. Burris, V., Smith, E., Strahm, A.: White Supremacist Networks on the Internet. *Sociological Focus* **33**(2) (2000) 215–235
3. Macskassy, S.: On the Study of Social Interactions in Twitter. In: Sixth International AAAI Conference on Weblogs and Social Media. ICWSM (2012)
4. 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
5. 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
6. Bermingham, A., Conway, M., McInerney, L., O'Hare, N., Smeaton, A.F.: Combining Social Network Analysis and Sentiment Analysis to Explore the Potential for Online Radicalisation. In: Proceedings of the 2009 International Conference on Advances in Social Network Analysis and Mining. ASONAM '09, Washington, DC, USA, IEEE Computer Society (2009) 231–236
7. Sureka, A., Kumaraguru, P.: Mining YouTube to Discover Extremist Videos. *Information Retrieval* (2010) 13–24
8. 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
9. Caiani, M., Wagemann, C.: Online Networks of the Italian and German Extreme Right. *Information, Communication & Society* **12**(1) (2009) 66–109
10. Caiani, M., Parenti, L.: The Spanish extreme right and the Internet. *Análise Social* **XLVI**(201) (2011) 719–740
11. Zuev, D.: The movement against illegal immigration: analysis of the central node in the Russian extreme-right movement. *Nations and Nationalism* **16**(2) (2010) 261–284
12. Blee, K.M., Creasap, K.A.: Conservative and Right-Wing Movements. *Annual Review of Sociology* **36** (2010) 269–286
13. McNamee, L., Peterson, B., Pena, J.: A Call to Educate, Participate, Invoke and Indict: Understanding the Communication of Online Hate Groups. *Communication Monographs* **77**(2) (2010) 257–280
14. Simi, P.: Why Study White Supremacist Terror? A Research Note. *Deviant Behaviour* **31**(3) (2010) 251–273
15. Greene, D., Reid, F., Sheridan, G., Cunningham, P.: Supporting the Curation of Twitter User Lists. In: 2nd Workshop on Computational Social Science and the Wisdom of Crowds at NIPS 2011. (2011)
16. Kwak, H., Lee, C., Park, H., Moon, S.: What is Twitter, a social network or a news media? In: Proceedings of the 19th international conference on World Wide Web. WWW '10, New York, NY, USA, ACM (2010) 591–600
17. 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)
18. Lancichinetti, A., Fortunato, S.: Consensus clustering in complex networks. *Sci. Rep.* **2** (03 2012)
19. Lancichinetti, A., Radicchi, F., Ramasco, J., Fortunato, S., Ben-Jacob, E.: Finding statistically significant communities in networks. *PLoS ONE* **6**(4) (2011) e18961
20. Hubert, L., Arabie, P.: Comparing partitions. *Journal of Classification* (1985) 193–218
21. Takhteyev, Y., Gruzd, A., Wellman, B.: Geography of Twitter networks. *Social Networks* **34**(1) (2012) 73 – 81
22. Kulshrestha, J., Kooti, F., Nikraves, A., Gummadi, K.: Geographic Dissection of the Twitter Network. In: Sixth International AAAI Conference on Weblogs and Social Media. ICWSM (2012)

---

2 <http://rossdawsonblog.com/weblog/archives/2012/02/which-countries-have-the-most-twitter-users-per-capita.html>