

Measuring Political Violence

Social Media Mapping



POLITICAL CAPITAL
POLICY RESEARCH & CONSULTING INSTITUTE





Co-funded by the Prevention of and Fight against Crime Programme of the European Union.

This project has been funded with support from the European Commission. This publication reflects the views only of the author, and the European Commission cannot be held responsible for any use which may be made of the information contained therein.



This project has been supported by the Foundation Open Society Institute.

MEASURING POLITICAL VIOLENCE

– Social Media Mapping –

October, 2015



POLITICAL CAPITAL
POLICY RESEARCH & CONSULTING INSTITUTE

CONTENTS

Why is the role of social media important in spreading extremist ideas	4
Methodology	4
United Kingdom Results	6
Hungarian Results	13
Strengths and weaknesses of these methodologies	20
Future uses	22

Why is the role of social media important in spreading extremist ideas

There has been a change over the last decade in the way people access, consume and produce media: a shift away from mainstream media and toward internet-based content and social media. Hateful or offensive content has been present on the internet from its inception. For example, *Stormfront*, recognised as one of the first 'hate sites', has operated since at least the early nineties. However some believe that social media has made it easier to publish and spread extreme or offensive views.

Radical right wing parties and movements are well established to be early and active users of social media, both as a way of producing cheap and rapid propaganda; creating a coherent group identity, and organising events and activities¹.

This study examines the way a selection of populist right wing pages on Facebook and Twitter accounts can shed some light on the way these groups use social media. Although it is increasingly recognised that these groups are active users of social media, there is a lack of research into precisely how they use it.

It is important to stress that both Facebook and Twitter have terms and conditions which prevent hateful and extremist groups and content being on the platform. For instance, Facebook prohibits and removes hate speech, which it defines as 'content that directly attacks people based on their: race, ethnicity, national origin, religious affiliation, sexual orientation, sex, gender, or gender identity, or serious disabilities or diseases'. Although they do not allow hate speech, sometimes people post disagreeable or disturbing content that does not violate the platforms' policies.

Methodology

This study was conducted using both Facebook and Twitter. The reason both sites were used was because both are known to be popular with the groups in question; and both platforms allow researchers to collect and analyse data from them in a relatively easy and structured manner.

It is possible to manually collect social media data in a number of ways - copying, screen grabbing, note-taking, and saving web-pages. However, where large volumes of data are involved, the most appropriate method is to collect the data automatically. This is done through connection to a platform's 'Application Programming Interface' ('API'). The API is a portal that acts as a technical gatekeeper of the data held by the social media platform. They allow an external computer system to communicate with and acquire information from the social media platform. Each API differs in the rules they set for this access: the type of data they allow researchers to access, the format they produce this data in, and the quantities that they produce it in.

Some APIs can deliver historical data stretching back months or years, whilst others only deliver very recent content.

¹ See Demos (2012) *The New Face of Digital Populism* for an overview.

Some deliver a random selection of social media data taken from the platform, whilst others deliver data that matches the queries – usually keywords selected by the analyst – stipulated by the researcher. In general, all APIs produce data in a consistent, ‘structured’ format, and in large quantities. Facebook and Twitter’s APIs also produce ‘meta-data’ – information about the data itself, including information about the user, their followers, and profile. This meta-data can be a rich source of information of value to social media researchers, often containing information on everything from the sender’s device type, to their account creation date, location and social media following².

There are several types of API access to Facebook data, most of which have been designed for app makers, such as Public Feed API, a Keyword Insights API, a Marketing API and Atlas API³. For this work we used the Public Feed API which allows researchers to access all data that has been posted on a public Facebook page. (Access to all Facebook data is predicated on the user’s settings and who has agreed to share information with them. Facebook’s privacy structures are complex – potentially, any single user can have a distinct privacy setting for every piece of data they share. The Public Feed API will only return data that is public).

Using the Public Feed API, we collected data from selected Facebook pages from the UK and from Hungary. We used ‘R’, an open source software that allows researchers to access publicly available data from public pages. These pages were handcrafted by the researchers who are subject matter specialists in the subject.

We also collected tweets via Twitter’s ‘stream’ and ‘search’ application programming interfaces (APIs). The ‘search’ API returns a collection of relevant Tweets from an index that extends up to roughly a week in the past. The stream API continually produces tweets that contain one of a number of keywords to the researcher, in real time as they are made. Identifying specific accounts also allows researchers to collect the last 3,200 tweets from that account. These tweets are then returned to the researcher’s own computer data set in a Json file, which can then be subject to analysis. We analysed the data using a software package called Qlik. As for Facebook, we subjected the data to a series of analysis.

In both instances we did not attempt to collect or use any personal information about individuals; nor did we attempt to identify any individuals. We have obscured all account names and quotes to ensure no individuals are identified, and deleted all the data following the analysis. (Although a decision was taken to mention some large organisational accounts). The purpose was to understand the broad patterns of behaviour. For Facebook we did not collect any data from groups or from individual’s pages; and we did not collect any data from closed pages. We did not attempt to collect or use any personal information about individuals; nor did we attempt to identify any individuals. Where a user’s name or ID was collected inadvertently, it was deleted.

2 C Miller, ‘Social Action on Social Media’, Nesta Working Paper, Working Paper Series (Nesta: 2015)

3 <https://developers.facebook.com/docs/graph-api/other-apis>

It is important to stress that these are in many cases quite experimental methodologies. There are no firmly established 'best practice' methods to collect and analyse data of this nature. Further, this is designed as a scoping study. Therefore findings need to be read with caution.

United Kingdom Results

Facebook

Over 2 months (1st October 2014 – 1st December 2014) we collected 497 posts from Facebook pages of ten populist right wing groups from the UK, which gave a total of 930,145 interactions (an interaction can be a 'like', 'share' or 'comment'). We tried to collect a fairly wide spread of different groups that have a presence on Facebook; and a number that was broadly comparable in size to the Hungary case study. Based on the experience of the researchers, we selected the following: Britain First⁴, British National Party⁵, English Defence League (EDL)⁶, EDL Forum⁷, English Nationalist Alliance⁸, I Am Proud to be British⁹, Infidels of Britain¹⁰, Send the SAS to Catch Jihadi John¹¹, South Wales British Movement¹², and the Yorkshire Angels¹³ (a female division of the EDL). However, there are many more similar Facebook pages that could have been part of this study – in particular, there are at least 20 EDL pages run by regional divisions.

It is important to stress that we do not claim that any of the pages included in this study are hateful, or that the content there is 'hateful'. Instead, we have focused this study on populist right wing Facebook pages, which are frequently accused of being a place where a high volume of hateful content is posted or shared. We refer to these pages as 'populist right wing' pages throughout, and have found that there is a very wide range of content posted and shared there. 'Posts' in this sense refer to updates that were made on the page by the administrator(s) of that page.

Overall data on the size and scale

On a measure of total 'page likes', the UK based Page Britain First is the most popular page of its kind in Europe: with 569 thousand page likes; and a total reach of 51 million users at the time of writing¹⁴. (This data is only available to page admins. It was publicly shared by the administrator of Britain First on their public page.)

4 <https://www.facebook.com/OfficialBritainFirst?fref=ts>

5 <https://www.facebook.com/OfficialBritishNationalParty?fref=ts>

6 <https://www.facebook.com/EDL-English-Defence-League-238696516197018/timeline/>

7 <https://www.facebook.com/English-Defence-League-EDL-Forums-130356320328490/timeline/>

8 <https://www.facebook.com/English-Nationalist-Alliance-170875669635671/timeline/>

9 <https://www.facebook.com/ImProud2BeBritish?fref=ts>

10 <https://www.facebook.com/INFIDELS-OF-BRITAIN-352629524764287/timeline/>

11 <https://www.facebook.com/CatchJihadiJohn?fref=ts>

12 <https://www.facebook.com/South-Wales-British-Movement-775558882477451/timeline/>

13 <https://www.facebook.com/Yorkshire-Angels-English-defence-league-168144953236977/timeline/>

14 At the end of September, 2015 number of page likes was almost 924 thousand.

On average, posts in our collection had 1871 interactions (which includes likes, comments and shares) per post. However, this is skewed by a small number of highly active posts. If we remove the top 10 most popular posts, the average is 1,200).

This masks enormous variation. The most popular piece of content had over 100 thousand interactions (below); and was posted by the group 'Britain First'. (See below). This is a photo of Prince Harry and William. Typically, Britain First posts content which is not overtly political in nature. By contrast, 30 per cent of the posts had fewer than 100 interactions. This illustrates the way in which a small number of very popular or 'viral' posts can have a dramatic reach (at least in terms of online reach), and a large proportion of posts have a relatively small reach.

The most popular Facebook post¹⁵



An analysis of network / membership structure and how that effects how content is shared

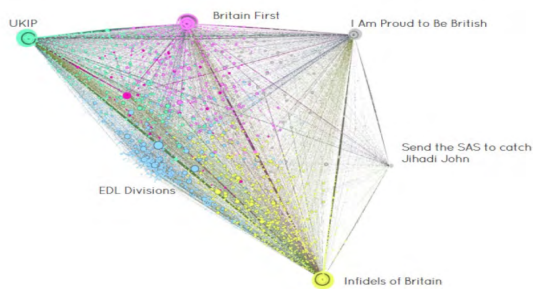
In order to better understand the way information and ideas flow across these pages, we selected the UK pages and examined the extent to which individuals who commented in one page also commented on another page. This data was taken from Facebook's API through R and visualized in Gephi, an open source network analysis tool.

¹⁵ <https://www.facebook.com/OfficialBritainFirst/photos/a.347167375428530.1073741829.300455573433044/573245729487359/?type=3&theater>

The images below show how the users are clustered around the different groups they are active in, including how active they are in each group they are members of. This is called 'edge rank'.

To build a network map, we extended our sample to 92 pages Facebook pages, which contained a total of 54,495 unique users, who between them made 159,437 comments on those pages. (This is calculated by collecting the comments and working out the number of unique user IDs that contributed to that data set). We extended the sample in order to make a more meaningful network map, as we wanted to examine how our pages fit into a wider network of users. To extend the network, we added a number of UK Independence Party (UKIP) Facebook pages; and added a number of EDL regional division pages. This analysis finds that 16.2 per cent of users are active on two or more pages; 1.3 per cent are active in four or more. It suggests that there are a small number of highly active users that create content in several pages.

Network of users




What type of posts are most effective at reaching a wide audience

In order to determine what type of posts were successful in reaching a wide audience, we examined a) the format of posts and b) the content and tone of posts.

Facebook API data allows researchers to determine what format posts take, divided by 'link', 'photo', 'status' or 'video'. The figure below shows what type of format was most widely used by the pages in question. Photos are extremely popular among these groups – which illustrates how important visual images are to these groups.

Content Types (rounded to nearest %)



Analysis of post content

In order to have a more nuanced understanding of what type and what tone of post is most popular, we analysed 129 of the most highly interacted with content from the 597 posts collected from these pages. We divided these into categories of content of post and tone of post. (The categories were selected by researchers based on a coding category created for this research. We judged post popularity by the average number of interactions they received).¹⁶

Content	Average of Total Interactions	Count of Category
Attack	27,302	2
Comment	28,005	48
News	9,978	60
Question	8,825	19
Grand Total	16,785	129

Tone	Average of Total Interactions	Count of Tone
Angry	14,857	61
Celebratory	19,451	30
Neutral	17,775	38
Grand Total	16,785	129

¹⁶ In some cases samples were too small to be confident in precise percentages, so it's best used as illustrative data.

The most popular tone of posts on populist right wing pages was 'celebratory' such as in posts commemorating war dead or patriotic pride (often involving an image)¹⁷. The most popular content of post is 'comment' which we define as a 'catch-all' label applied to content that is created or shared by the page that commentates on a situation without necessarily referencing outside sources. These are very high averages because they are often skewed by a small number of very highly shared content.

On these pages, 5 per cent of all comments were categorised as counter speech, meaning comments which disagreed with the post or presented an alternative, more positive message.

Twitter

Researchers created a handcraft set of 30 'seed accounts', based, as far as possible, on a similar selection of groups as the Facebook pages. These were public accounts which were known to the researchers to be active and open supporters of the English Defence League (a UK street based anti-Islamism movement) or the British National Party. Based on these seed accounts we then built a network out of users who followed those accounts and then collected tweets posted by those users in the past 3 months. We have anonymised all accounts and not named any individual user or included any posts, to protect user privacy. Rather, the purpose of the work was to better understand the nature of the network of users.

Structure of network

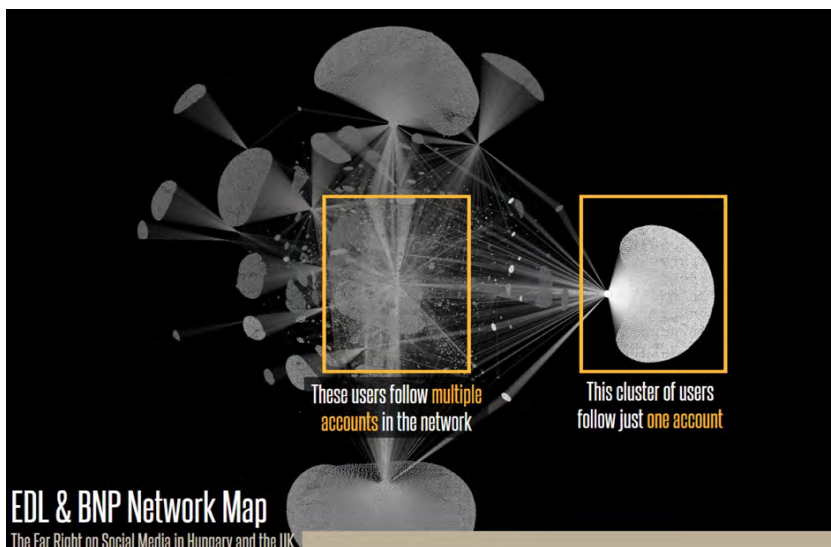
By building up a network in the way described above, we were able to determine the general shape of the network of accounts associated with right wing populist accounts on Twitter.

Based on our analysis, we found a total of 46 thousand users in the network. Of those, 8 thousand follow at least two other users; around 3,700 follow at least 3, and around 1,500 follow at least five. Two-hundred and ninety six – which might be considered 'the hardcore' network – follow at least ten.

We took the 46 thousand accounts and produced a network map to illustrate the extent to which different users follow multiple accounts within the network. These maps are built using Gephi, an open source network analysis tool. We plotted each user and page on the map. If a user had interacted with a post, a connection (or 'edge') was drawn between the two users.

As the diagram below illustrates, there are small clusters of activity which are not always embedded in the entire network. The cluster to the right in the diagram is of British National Party supporters – which highlights the fact these users follow just a single account from the network, and not others users.

17 Celebratory content is that which celebrates the page or its values.



Levels of activity

Of the 1,500 'core' accounts identified, 1,040 had been active in the last three months (around 69 per cent of the total). We then collected data from those accounts, and found that – over the last three months – they had posted 795 thousands tweets. This averages at around 760 tweets per account over the period in question, which is around 8 tweets per day per account.


In order to have a better understanding the sort of content being posted by these accounts, we collected the top 'hashtags' being used in tweets¹⁸. The chart below shows the top 15 hashtags used over a three month period – and the number of times they were used. This shows a number of things.

- First, there is a lot of engagement with (or discussion about) mainstream media outlets and stories. UKIP, the BBC and BBC Question time are all popular hashtags among the network and all relate in some sense to mainstream subjects.
- Second, as expected, there is a considerable amount of discussion relating to Islam and Islamism.

¹⁸ A hashtag is a typically just a word with an accompanying # sign, and is used to allow the content to be more easily findable by other users who search Twitter for those hashtags.

- Finally, the network has used hashtags that are associated with right wing (or extreme right wing) politics. 'Tcot' refers to 'Top Conservatives On Twitter' and is typically used in the United States. More extreme, 'WhiteGenocide' is a phrase sometimes used by extreme right wing groups to describe what they consider a general diminution of the white race around the world – typically as part of a wide plot by various shadowy groups.

Top Hashtags



Across the network overall, we searched for the most popular single tweet. It related to a call by one user to ban the burka in the UK. It included an image and encouraged other users to re-tweet it, which was done 3,815 times. Encouraging someone to re-tweet content is a common strategy to increase its reach. It is not possible however to determine how many other users may have seen it – although it is likely to be well in excess of one hundred thousand people.

Hungarian Results

Facebook

Between February 2014 and June 2015 we collected 13,877 posts from 18 public Hungarian Facebook pages, which gave a total of 3,501,055 interactions (an interaction can be a 'like', 'share' or 'comment'). Most of the selected Facebook pages were connected to far-right groups or movements; some support pro-Russian propaganda; and a small number could be described as militaristic or spreading conspiracy theories. These pages were selected by researchers based on their knowledge of the subject (but they are not all named, in order to protect individual privacy – only the largest pages are mentioned by name). As above, we did not collect any information from private or password protected pages. As for the UK research, it is important to stress that we do not claim that any of the pages included in this study are hateful, or that the content there is 'hateful'. Instead, we have focused the Hungarian research on far-right wing, militaristic or conspiracist Facebook pages, which are frequently accused of being a place where a high volume of hateful content is posted or shared.

Overall data on the size and scale

Within our set of Facebook pages Jobbik's official page is the most popular, with 298,876 likes as of end of September 2015.

On average, posts in our collection had 252 interactions (which includes likes, comments and shares) per post. This is by far less, than the average interaction in the UK sample (1,871). The most popular post appeared on the page of Jobbik and had 40,842 total interactions (29,134 likes, 732 comments and 10,976 shares). This was a celebratory photo posted to commemorate international women's day, showing young women at a wedding ceremony in traditional Hungarian clothes.

The most popular Facebook post¹⁹



The number of interactions generated by the second most popular post was less than half of this. The top five posts appeared on the Facebook page of Jobbik, but only two of them were celebratory: one was critical of the government (against the government's planned Internet tax), one was a nationalistic (supporting the use of buses that were manufactured in Hungary) one was critical of the police, and the other celebratory posts celebrated Jobbik's first winning candidate in a single-member constituency.

¹⁹ <https://www.facebook.com/JobbikMagyarorszagertMozgalom/photos/a.10152008068391405.1073741874.287770891404/10153043631511405/?type=1>

The 2nd – 6th most popular Facebook posts



**A FELVIDÉKI NEM SZLOVÁK!
A KÁRPÁTALJAI NEM UKRÁN!
A DÉLVIDÉKI NEM SZERB!
AZ ERDÉLYI NEM ROMÁN!
ŐK IS MAGYAROK!
EGYETÉRTESZ? OSZD MEG!**

www.fb.com/JobbraAt


As Jobbik's Facebook page is the most liked one in our sample, it is not surprising that among the first 85 most popular posts there is only one²⁰ which was posted on another page (Jobbra Át, or 'Turn to the right'). This was an image saying that ethnic Hungarians living in neighbouring countries (that were part of Greater-Hungary before Trianon) are also part of the Hungarian Nation. This post was the 6th most popular with 13,292 interactions (most of them, 11,768 were 'shares').

²⁰ <https://www.facebook.com/JobbraAt/photos/a.201859210010910.1073741828.201838506679647/349345621928934/?type=1>

The 100 most popular posts (ranked by the number of interactions they have generated) are dominated by photos (89 of them were photos, 10 were links and there was 1 video).

It is important to mention that a post that linked to an article on Alfahír.hu (an online media outlet connected to Jobbik) generated 10,551 interactions. The article²¹ reported about a crime committed by gypsies against young non-gypsy Hungarians. This high level of popularity reflects how anti-Roma attitudes are present and embedded in the Hungarian society and in particular among sympathizers of Jobbik²².

An analysis of network

In order to better understand the way information and ideas flow across the selected Hungarian Facebook pages, we examined the extent to which individuals who directly engaged with a page by commenting on one page directly engaged with another page. This data was taken from Facebook's API through R and visualized in Gephi, an open source network analysis tool.

Our sample contained a total of 39,514 unique users, who made 168,334 comments on those pages. (This is calculated by collecting the comments and working out the number of unique user IDs that contributed to that data set).

Most of the commenters are active on only one page (83 per cent). A further 12 per cent are active on two pages and slightly more than 5 per cent commented on at least 3 fanpages. Regarding number of comments, the most popular fanpage is Jobbik's, with 124,303 comments by 30,691 unique users on the last 1,000 posts between 13th October 2013 and 24th July 2015. The second most popular page is Alfahír, where 6,534 unique users made 29,265 comments on the last 1,000 posts. (These 1,000 posts were posted since April 2015 meaning that while Jobbik's community is much larger, the community engaging with Alfahír is more active). In both cases these comments then spark a conversation in the replies to that initial comment.

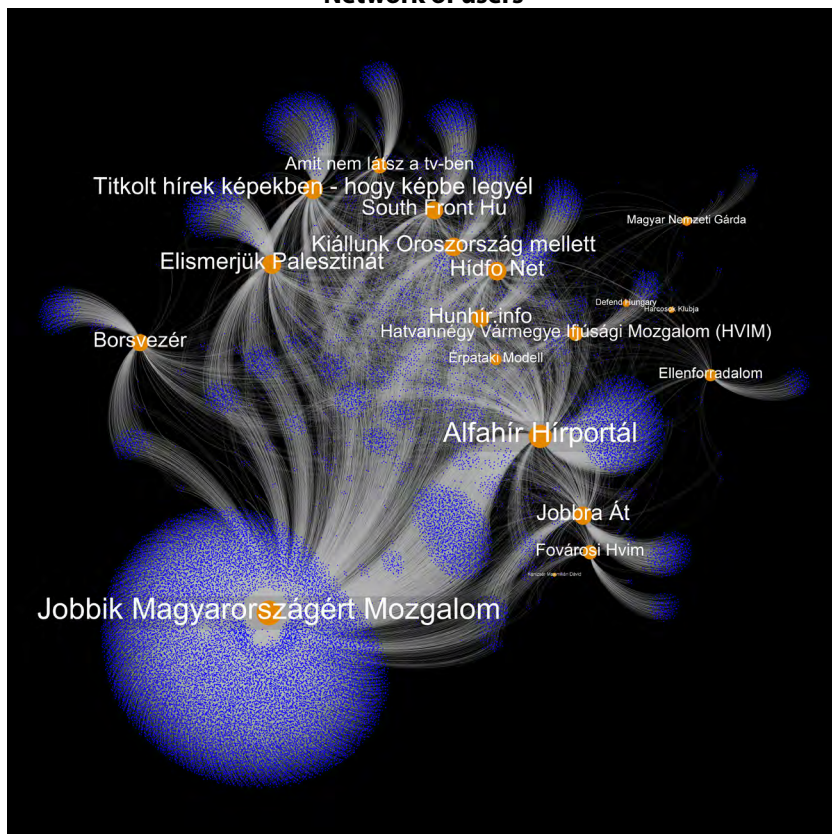
More than half of the users active on all pages commented on only one post (52.75%). 15.6 per cent commented twice. 84 per cent of the active users have no more than 5 comments. Only 8 per cent commented on more than 10 times.

The graphic below shows how the users are clustered around the different groups they are active in, including how active they are in each group they are members of. This is called 'edge rank'. Clouds of blue dots that are connected to only one Facebook page (coloured in orange) represent users who commented on only one page. Clouds that are connected to more than one page contain users who are active on more than one Facebook page.

21 http://alfahir.hu/ugy_megverte_a_cigany_gyerek_az_iskolaban_hogy_koponyatorese lett

22 Bernát, Juhász, Krekó, Molnár: *The roots of radicalism and anti-Roma attitudes on the far right*, 2013. Available online: http://www.tarki.hu/en/news/2013/items/20130305_bernat_juhasz_kreko_molnar.pdf

Network of users



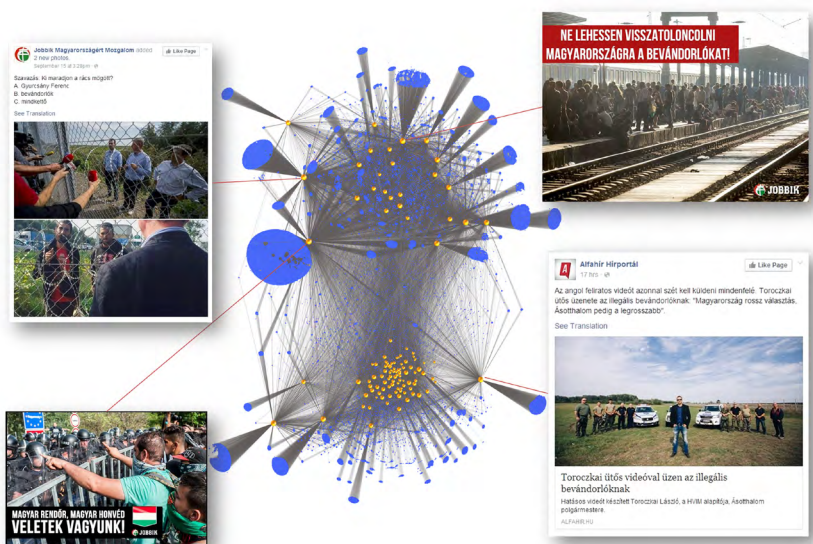
As noted above, Jobbik's page received the most comments by far: but the majority of these commenters are active on only this page. On the upper half of the image the group pro-Russian pages can be seen close to each other. This means that a large part of their commenters are active on 2 or more of them. Conspiracist pages tend to be relatively distinct – on the top of the image, but through their common commenters they are also connected to the network.

The refugee crisis on the Facebook page of Jobbik and Alfahír

Jobbik has an extremely anti-migrant (pro-nationalistic) stance in the current refugee crisis. This is clear from their Facebook page – most of their posts were connected to this issue. To get a detailed view on how these posts spread, we conducted further research of the pages of Jobbik and Alfahír between 12-17 September, 2015. We took the 118 posts that appeared during this period and produced a network map to illustrate the extent to which different users interacted with these posts.

We plotted each post on the map with yellow dots. Unique users, who either liked, commented or shared these posts are represented by the blue dots. If a user had interacted on two particular posts, a connection (or 'edge') was drawn between the two points. Users with a high number of edges had interacted with a number of different posts.

As the diagram below illustrates, the posts are clustered into 2 separate groups. In the upper group are the posts appeared on the Facebook page of Jobbik, while in the lower groups are posts on the Alfahír page. We can see some edges connecting these groups, which means there were users who interacted with posts from both of these pages: 8 per cent of users liked posts from both Jobbik and Alfahír, while 73 per cent liked only Jobbik's posts and 19 per cent liked Alfahír's posts. This means Jobbik can reach to a much wider audience via its Facebook page than Alfahír can. Some posts (highlighted on the image) were very popular and a lot of users interacted with them, who generally are not active on these pages (the large blue spots with only one edge). This means that these images could trigger a lot of attention from Facebook users who are not core sympathizers of Jobbik.




Twitter

Researchers created a handcraft set of 24 Twitter 'seed accounts' from Hungary. These were accounts which were known to the researchers to be active and open supporters of Jobbik (regional, local organizations of the party). Based on these seed accounts we then built a network out of users who followed those accounts and then collected tweets posted by those users in the past 3 months.

Structure of network

We produced a network map to illustrate the extent to which different users follow multiple accounts within the network. As the figure below illustrates, similarly to the UK result, there are small clusters of activity which are not always embedded in the entire network. The cluster in the lower-right corner represents supporters of Jobbik's official Twitter account – which highlights the fact these users follow just a single account from the network, and not others users, just like BNP supporters in the United Kingdom.



Jobbik Network Map

The Far Right on Social Media in Hungary and the UK

Levels of activity

We identified 499 'core' accounts, these followed at least four others in the network. Among them only 127 unique users had been active in the last three months (around 25 per cent of the total, a far less number than we got in the UK research). We then collected data from those accounts, and found that – over the last three months – they had posted 29,469 tweets. This averages at around 232 tweets per account over the period in question, which is less than one tweet per day per account. This is also significantly less active than the groups studied in the UK.

Strengths and weaknesses of these methodologies

There is a growing interest in the field of 'big data' analytics and how it can be applied for social science research – including for the study of radical groups and movements. There is a considerable amount of relevant and useful data available on social media platforms, particularly Twitter and Facebook (although there are many others). Typically, social media datasets are far larger than comparative datasets gathered through conventional polling, interviewing and surveying techniques. This often (although not always) means it requires automated systems to collect and analyse them and this creates new methodological challenges. Social media data are often – although not always – generated by users themselves, and sometimes contain personal data, and this also creates new ethical concerns.

Overall, while there has been a considerable growth in the type and nature of social media research – especially in the commercial and advertising sectors – there is no broadly accepted best practice body of methodology and ethics about how to collect, use, and present these data sets for academic and social science researchers. Established ways of researching attitudes have long histories of use. This experience has consolidated into a body of good practice – dos and don'ts – that, when followed, ensures the quality of the research. Social media research does not have a long history of use, or a collective memory of what works and what doesn't. It uses new technologies in ways that are unfamiliar with the social sciences, often with new and important implications for research.

This research has found that:

- Social media has become an important and active venue for a wide variety of populist right wing activity.
- However, although the network appears large, it is often led by a relatively small number of dedicated and active users.
- Nevertheless, these pieces of content can reach relatively large audiences outside their own network of users.
- They use both Facebook and Twitter extensively to discuss and share information about both mainstream and very niche political issues. It is common for these groups to share 'mainstream' information in order to reach a wider audience.
- For reasons discussed below, however, we are very hesitant to draw general conclusions based on these data sets about the offline composition of these groups.

In very general terms, most social media data tends to share the following broad features, which are useful in terms of deciding when and where it is (and where it is not) a potentially useful source. These are both positive and negative attributes of these data sets.

Strengths

Relational: Because most social media is premised on curated networks of users, most data include some information about the relationship between users. This can take several forms: for example, if a user follows another user; has posted to another user; has interacted with another user; or has shared another user's content. What these relationships mean remains an open research question.

Real or near-real time: Many social media platforms allow data to be collected as soon as it is posted. For example, on Twitter, researchers can access tweets as they are posted by users, making real time research work possible.

High volume, available at low (or even no) cost: One of the major benefits of this type of research work is that open social media data is available, often for free via APIs, and at very large scale.

A new way in to understanding these groups: Social media is an increasingly important way in which various groups – including radical groups – communicate with each other, share content, and build a sense of identity. It is, therefore, also a space that researchers need to understand in order to build up a better understanding of how they operate.

Reactive and indirect: Social media is often a reactive source of data; a space where people react to an event – either online or offline. This creates a dynamic relationship between media reports and stories and broader conversations which take place afterward; and creates new challenges in respect of accurately determining opinions and attitudes, which are often indirectly expressed.

Weaknesses

Demographic and self-selection biases: many social media users do not demographically represent wider populations (they remain slightly younger, and more urban than average). Anecdotal and small-scale research suggests they might also be more liberal than average. Moreover, even collected data often does not represent all users, because it appears that many users go on Twitter or Facebook to express a particular reaction to an event if they have a strong opinion about it, and so they are not necessarily a representative sample even within the platform.

Unpredictable: It can be extremely difficult to predict in advance the likely volume and data quality of social media data on any given subject. This can make it difficult to plan in advance what topics and subjects can be researched.

Forum specific biases: Social media spaces are new social spaces, which are characterised by their own norms and mores. For example, based on our research, Twitter is a medium characterised by humour, sharing stories, and anti-establishment sentiment. For a human analyst not habituated with certain memes or group specific language it can be very difficult to determine likely sentiment or underlying attitude. This is even more difficult, if not impossible, when training any automated system to recognise these very subtle distinctions.

Future uses

In summary, we believe, if used with careful methods and caveats, this type of research can be very useful to better understanding the nature and beliefs of groups – providing they have a presence on social media. In particular, we suggest it is of most value for the following purposes:

Understanding trends in thinking: Social media research offers a unique opportunity to understanding trends in thinking and beliefs within a group or set of groups. For researchers interested in how beliefs evolve over time – including in more or less violent directions – this type of research is invaluable.

Group response to events or flashpoints: Social media is a reactive platform. This means that much online traffic on both Twitter and Facebook tends to be driven by recent events, and individuals' response to them. This also provides insight into the way a group responds to specific external stimuli, which in turn can help provide greater understanding to how certain events are likely to provoke a response.

Measure size and reach of content: Most social media data includes 'meta-data' such as volume of interactions, re-tweets or shares. This allows researchers to gauge the possible reach of certain pieces of content – and indeed what is popular content among users. This could be 'engagement' (for example, it might be the ratio of users who viewed the page and those who signed up, which is a useful proxy of potential reach. We have been able to measure this effectively; volume and exposure, such as how many posts are being produced on the topic, and how many unique users are discussing it? Or how large is the audience? Is hate speech being limited to isolated communities (either by the communities themselves or Facebook's personalisation algorithms)?

Understanding networks: Network maps are relatively easy to construct, and provide a useful illustration of influential accounts or users within a data set – either to measure one's own position or to identify other important stakeholders talking on a subject.

Counter-speech measures: While this report has focused on populist right wing groups; the same approach can be applied to better understand users who are confronting or disagreeing with certain groups or ideas. This can be used to identify areas where the reach and engagement of this content could be increased.

On the whole, we are sceptical about these approaches as a way to 'predict' or 'spot' violent intention or behaviour. It is better understood as a useful research instrument for academic (and other) researchers, used in conjunction with other techniques. However, there are a number of challenges involved with using this research in a way that is useful for researchers, academics, and public policy specialists. Because it is a new discipline there are several considerations that should be met when decided whether to undertake this type of research or not.

Does the research question require social media data – and does it need to be automated? Based on the features of social media data as set out above, not all research questions require, or would benefit from, a social media component. A compelling case should be made as to why social media is a valid instrument of study for the research question, and why it is preferable to existing, more established research methods.

Clarity over data access: So called ‘black box’ data – where data is provided by a third party without clarity over methods, search terms used, or access levels – should be avoided wherever possible. This means that ‘off the shelf’ data analytics tools are likely to be less valuable than systems that allow researchers and analysts control over how the system operates.

Clarity over sampling methods: Data is sometimes acquired on social media platforms by something called key word matching. This is where data sets are trawled to identify and collect matches with (a) pre-determined word(s) or term(s). Because data are collected based on conversations rather than demographic or what we call ‘topographic’ details (for example, the power law features), it adds a high degree of uncertainty into the demographic background of any collected data set. These keywords can produce different kinds of problems – sometimes they are over-inclusive (and collect irrelevant data), and sometimes they are under-inclusive (and miss relevant data). In both these ways, key-word matching is inherently prone to systemic bias – meaning that the data collected and therefore the conclusions drawn, are affected in a non-random way by the search terms employed. With respect to data quality overall, it is important to make a distinction between internal and external validity. At present, for example, we do not believe Twitter is a valid instrument to conduct reliable, population-level opinion surveys. There are often significant problems with several types of self-selection bias in social media data and often no clear way to correct them. Statements making generalisations about overall public attitudes based on social media data sets – ‘external validity’ – should be made with extreme caution.

Adherence to research ethics: Conducting research using Twitter or Facebook data presents new ethical challenges in respect of how researchers should collect, store, analyse and present data. Because it is a new field of research, there are no widely accepted protocols and approaches for how to do this ethically. In the UK, the Economic and Social Research Council (ESRC) principles of ethical research²³ is an excellent guide for conducting research of all kinds – and can be usefully applied to online as well as offline research. Social media research should adhere to the research ethics standards set out by the ESRC’s principles. The key questions are whether or not the research has sufficiently explained the risks and minimisation strategies for:

- The potential identification of individuals.
- Whether or not the research has sought informed consent, and, if not, why it is not considered necessary (ideally with reference to the expectation of privacy a research subject might have).

²³ <http://www.esrc.ac.uk/funding/guidance-for-applicants/research-ethics/>

- Whether there is any possible harm to the individual, and what measures there are to minimise them.
- Whether techniques to 'cloak' or protect the identity of research subjects are necessary, and how that might adversely affect the quality of the research.

As a very general principle, where an individual is identifiable, explicit permission should be sought, unless a) it is clear that the subject has no expectation of privacy and b) the research will be significantly adversely affected unless the individual is identified.