Social Network Analysis Visualisation Report
Textgain Internship

Boris Marinov

February 2021

1 Introduction

The use of Twitter and other forms of social media platforms as "news-outlets" is well documented, with 68% of people in the United States reporting some form of reliance on these platforms for staying up to date with current events (Matsa & Shearer, 2018). This concerns not only source of health and science information (Hitlin & Olmstead, 2018), but also a spread of information in regards to previously large scale and rapidly occurring events (Househ, 2016; LaLone et al., 2017; Daughton & Paul, 2019). The platform's use in monitoring the public perception of the most recent COVID-19 pandemic has already seen some attention (Ordun, Purushotham, & Raff, 2020) with effective and insightful findings.

While a lot of positive and useful information is being shared, for example by credible organizations such as WHO, an emergent term of the "infodemic" outlines many of the serious problems of the platform. As there is no real gatekeeping of content, the spread of potentially false and harmful information is vast and misinformation detection is a long withstanding, continuously researched field (Guo, Ding, Yao, Liang, & Yu, 2019). Techniques for detecting misinformation largely fall under two categories: content and lexical analysis, as well as network analysis.

Social network analysis (SNA) and user activity is evidently an important source of information. The social networks being analysed often contain large collections of users with complex interactions, and detecting the flow of information is a challenging task. As the network is a form of graph, graph theory measurements such as centrality, degree and betweenness have already seen wide use in capturing the patterns of the information flow (Borgatti, 2005; Burt et al., 2005). Efforts to rank actors based on graph-based measures within clusters, in order to focus analysis on a much smaller subset, have also been undertaken (Ediger et al., 2010). Other research has focused on determining the speed and extend of Tsunami warnings within these networks (Chatfield & Brajawidagda, 2012) or more complicated methods which identify different types of information flow and use these to topically segment their networks (Himelboim, Smith, Rainie, Shneiderman, & Espina, 2017). SNA thus proves to be a useful tool in regards to large scale events, emergencies or newly emerging societal movements.

One such movement is the *QAnon* far-right conspiracy theory which started in 2017. The conspiracy theory states that the former US President, Donald Trump, is carrying out a secret war against high-level Satan-worshipping paedophiles in governments, businesses and media[1]. Despite being debunked, the idea has generated a large following. Combined with the recent US elections, political connections and the spread of misinformation, the QAnon movement is a suitable and inviting candidate for SNA application.

Textgain[2], a Belgian text-analytics company, has scraped a large collection of tweets related to the QAnon movement between the months of October and November (the time period leading up to the US election), with the intent of identifying the spread of the conspiracy theories into other parts of the world[3]. The following project aim is to apply SNA methods to the collected data. In particular, we aim at answering the following research questions:

- **RQ1:** Can we use the Twitter mentions in the collected data to create a 3D social network visualisation related to the QAnon movement?
- **RQ2:** Can NLP-based attributes extracted from the tweets be used to identify clusters of similar communities and thus influential users?
- **RQ3:** Can we visualise the flow, spread and interaction of the topics and locations in our dataset?

To create useful 3D visualisations we take steps in subsetting the data and using keywords as well as graph-based measures to focus on particular sets of users (**RQ1**). We make use of the locations of users from the data, topic modeling and Textgain's own AI tools to extract the NLP-attributes we want to look at and incorporate that into the graph to form clusters and single out influential users (**RQ2**). Finally, we attempt to collapse the graphs based on the different attributes to get a more general idea of the information flow (**RQ3**).

2 Methods

The following section will quickly go over the dataset details, how the data is modified for the visualisations, as well as details on the different attribute synthesis and main steps taken in the process of creating the graphs.

2.1 Dataset

Three main data files were created and provided by Textgain in regards to their investigation into the QAnon movement. The main data file contains over 0.6M

[1] `https://en.wikipedia.org/wiki/QAnon`

[2] `https://www.textgain.com/`

[3] `https://www.textgain.com/portfolio/qanon-spreading-conspiracy-theories-on-twitter/`

Tweets scraped with the use of predefined keywords (see Appendix). This file also contains information on the nature of the tweet, whether it is a reply or a retweet. Following, the users data file contains more detailed information on the users contained in the large tweet dataset, such as their profile description, location and number of followers and tweets sent since profile creation. The final mentions data file contains all the user mention relations found in the data, extracted from the complete dataset.

Currently the mentions file contains over 0.85M relations, which is too much to get a useful and clean visualisation out of. We can reduce this by focusing on users who have stronger affiliations and opinions to the QAnon movement. Using the keywords from the full dataset (see Appendix), we only keep the users whose username, screen-name or user profile description contains one of the keywords/phrases. This method already reduces the number of mention relations to around 120k pairs.

In addition to having the mention pairs, we are also interested in knowing what type of relation this is (reply or retweet). When building the graphs, knowing which node is the Target and which is the Source is already sufficient for basic visualisation. However, the addition of the NLP-based attributes and user information, needs to be done at the correct node. For example, in the case of a retweet, we want to add the attributes like polarity to both the Source and Target nodes, whereas in the case of a reply the polarity would only be added to the Target node. The type of mention relation is extracted by going over the complete dataset and checking if the keyword "RT" is contained at the start of the tweet.

2.2 Attribute Synthesis

In order to see whether the visualisations can provide and uncover certain patterns in the data, several different attributes are added to the mentions pair file before starting the process of building the graph. Ideally, nodes can later be coloured or grouped based on the attribute value or neighboring node attributes.

2.2.1 Tweet Polarity

Using Textgain's library `Grasp`[4], a polarity score can be calculated for the tweets associated with the mention pair subset found in the earlier stage. The overall distribution of the tweet polarities is as follows:

- 70809 neutral (polarity = 0) tweets (56.77%)
- 33695 positive (polarity > 0) tweets (27.01%)
- 20016 negative (polarity < 0) tweets (16.05%)

[4]`https://github.com/textgain/grasp`

2.2.2 Topic Modeling

Topic modeling is performed on the entire dataset to find cluster of words which describe individual topics/discourses. Associating a tweet with a specific topic would hopefully allow to further subset the data, as well as form more informative visualisations and find the clusters of interest we are after.

Short-text topic modeling is a challenge of its own. This is largely due to the sparsity of the formed matrices for most methods. LDA is a very popular topic modeling method (Jelodar et al., 2019), however it assumes that a text is made up of several topics (Shi, Kang, Choo, & Reddy, 2018). Tweets are however inherently short, and designed to reflect a single piece of thought from a user.

A recently proposed method called SeaNMF (Semantics-assisted Non-negative Matrix Factorization) (Shi et al., 2018) is applied to all tweets in the dataset. Existing NMF methods learn topics by decomposing the term-document matrix into lower ranked matrices, demonstrating strong performances in dimension reductions and clustering for high-dimensional data (Choo, Lee, Reddy, & Park, 2015), with the approach being successfully applied to topic modeling (Kim, Choo, Kim, Reddy, & Park, 2015). SeaNMF builds on top of this by leveraging word-context semantic correlations during training, overcoming the sparse problems of short texts, and thus outperforming LDA. Prior to application, some preprocessing steps are taken in order to get the most of out of the topic modeling (i.e remove words and phrases which are unlikely to carry meaningful information to describe a topic):

- Removal of hyperlinks
- Removal of punctuation, emojis and miscellaneous characters
- Removal of mentions
- Removal of closed-class words (pronouns, determiners, conjunctions, prepositions, punctuation and also numerals). Again, this is done using Textgain's Grasp AI toolkit.
- Removal of words shorter than 3 characters

An important parameter for the method is the number of topics. Table 1 evaluates the results of running the method with a different number of topics, using four separate measures. The Pointwise Mutual Information (PMI) measures indicate how well the words in each topic relate to each other (relying on co-occurrence), while the Topic Diversity (TD) and Rank-biased overlap (RBO) scores measure the diversity of the topics (Bianchi, Terragni, & Hovy, 2020). Based on a combination of the presented scores, we believed that 40 topics lead to the best and most interpretable results.

Typically stop-word removal is also performed. In this case however getting a stop-word list for all the languages included in the data might not be ideal or possible, so a custom list is created after initial topic modeling results. Certain

Table 1: Topic modeling evaluation scores

Num Topics	PMI	NPMI	TD	RBO
30	3.378	0.302	0.97	0.999
35	3.511	0.314	0.97	0.999
40 (stops removed)	**3.716**	**0.30**	**0.98**	**0.999**
40	3.467	0.321	0.96	0.998
45	3.494	0.321	0.95	0.998
50	3.570	0.286	0.95	0.998
65	3.830	0.310	0.95	0.998

topics end up being clusters of stop-words from different languages and the method is re-run with these words removed (running for 40 topics again).

A difficult part of any topic modeling procedure is the interpretation of the word clusters (topics). This is rather subjective, however an important part of the process as each message needs to be assigned to some topic. The full table of all topic words and assigned labels can be seen in the Appendix (Table **??**). Certain topics contain similar words and are grouped together, and topics which are deemed as irrelevant or too hard to interpret are labeled as **MISC**. Using this, each tweet gets a topic label which can then be incorporated into the mention relations and added as an attribute to the main graphs visualisations on.

2.2.3 User Location

Finally, we use the users data file to extract the location of the users in our network. This is done with the hopes of being able to visualise the outreach of the information being shared. We make use of the Python library `Geopy`[5] which is able to provide a country based on a location string. Twitter data is however far from standardized, as users can input whatever location they choose. In cases where `Geopy` was unable to retrieve a country, manual mapping is undertaken to be able to capture as much of the included locations as possible. For example, locations such as “The Great State of Texas” are mapped to “USA”. Other, less clear cases such as “UK/Poland” are left as “Not Found”. Following mapping, we are left with 42 Source user locations, and 142 Target user locations. This is already an indication of the wide global spread of the *QAnon* movement. Table 2 includes a summary of the top 10 Source and Target user locations discovered in our final resulting mentions file (“Not Found” locations not included).

2.3 Creating the Graph

Setting up the graph nodes, edges and exporting the file to a format which can be used by other software for visualising is performed with `Networkx`[6], a graph-based library in Python.

[5] `https://geopy.readthedocs.io/en/stable/`

[6] `https://networkx.org/`

Table 2: Percentage of the top 10 countries for both Source and Target users ("Not Found" not included in calculations)

Source Country	%	**Target Country**	%
USA	59.15	USA	45.65
France	14.07	France	13.91
Canada	7.31	Netherlands	12.55
Germany	6.72	UK	7.31
Brazil	4.73	Spain	5.61
Poland	2.09	Brazil	4.51
Luxembourg	1.98	Germany	4.4
Argentina	1.75	Israel	2.91
UK	1.17	Canada	1.91
Netherlands	1.03	Colombia	1.24

To do this, we first iterate over the mentions file and set up the Target and Source nodes based on the username, as well as the edges between them. At this stage the attributes (polarity, topic, user location etc) are added to the appropriate node, based on the relation type (retweet or reply).

We then check how many components there are in the formed graph. In graph theory, a component is a subgraph in which any two vertices are connected by an edge, and contain no additional connections to the rest of the graph. For visualisation purposes only the largest component is selected from the mentions file (there are many mention pairs which are not connected to the rest of the graph).

The next step is to go through the graph component and sort all nodes depending on their node degree. The node degree summarises how many incoming and outgoing connections each node has. Nodes with the highest degrees are likely cluster centers, and popular users in our network. Therefore, we select the top 40 users and subset the mentions file again to include pairs containing these users as either Sources or Targets. This reduces our final data to approximately 28k mention pairs.

We repeat the process of forming the graph and selecting the largest component again, and export the graph and all added information to a `graphml` file which can be used by the subsequent visualising method we select. At the same time, we use the built in layout functions from `Networkx` to position the nodes in our network with a force-directed algorithm and store 3D coordinates for each in a dictionary which later be used for plotting.

3 Graph Visualisations

3.1 Selecting Graphing Method

SNA visualisations are typically done in a 2D fashion. We instead aim for a 3D approach with the hopes of increasing interpretation of our results, as well as allowing for more complex patterns and relationships to emerge. Finding a 3D visualising software that can efficiently plot the size of our network, as well as provide nice functionalities which will help address our research questions proved to be trickier than anticipated. Initially we hoped to make use of the popular network visualising software Gephi[7]. Gephi has however been depreciated, and the previously existing 3D plugins (ForceAtlas 3D) are no longer available for use.

Our second idea was to create all plots manually since we have access to the individual nodes' coordinates. Matplotlib offers the option of creating 3D scatter plots[8], and an initial attempt can be seen in Figure 1. While we are able to create the visualisations, the limitations of using this method became apparent quickly. Firstly, navigating and moving around the plot does not offer much in return. One can rotate the plot, however selecting individual nodes and storing some information in them that can be accessed interactively is impossible. Furthermore, any specific node colours, changes or running specific algorithms on the graph need to be done manually and this introduces a lot of room for error, as well as time and computational constraints.

Instead, we make use of a relatively new and continuously updated visual analytic application called Graphia[9]. The benefits of using Graphia are several. It provides an interactive graphical interface where the exploration of the network can easily be done, even on less powerful computers. A user can easily find specific nodes they are after, whether it is based on a node name or a commonly shared attribute. It also has many built-in algorithms and methods which can be used to create the force-direction graphs, reduce edges and nodes, filter out components or collapse nodes based on a common attribute and so on. Lastly, it offers good editing and export support. In short, it provides us with a lot more flexibility, reliability and speed in addressing our research topics.

3.2 Initial Graphs

Figure 2 shows the complete network visualisation of all users selected based on the keywords only. Figure 3 shows the resulting visualisations with the keyword selection procedure, as well as the selection of top 40 node degree users as mentioned earlier. Markov Clustering (MCL)[10], a widely used unsupervised clustering algorithm in the field of bioinformatics, is performed using Graphia in order to assign colours to groups of nodes. While visible clusters can be observed in Figure 2, the comparison between the two visualisations shows the need to subset the mentions relations even further. Figure 3 shows a cleaner, less cluttered plot where clusters are more evident and we are unlikely to lose a lot of information from many individual nodes with scarce and weak connections.

---

[7] https://gephi.org

[8] https://matplotlib.org/api/_as_gen/mpl_toolkits.mplot3d.axes3d.Axes3D.html?highlight=scatter3d#mpl_toolkits.mplot3d.axes3d.Axes3D.scatter3D

[9] https://graphia.app

[10] https://micans.org/mcl/

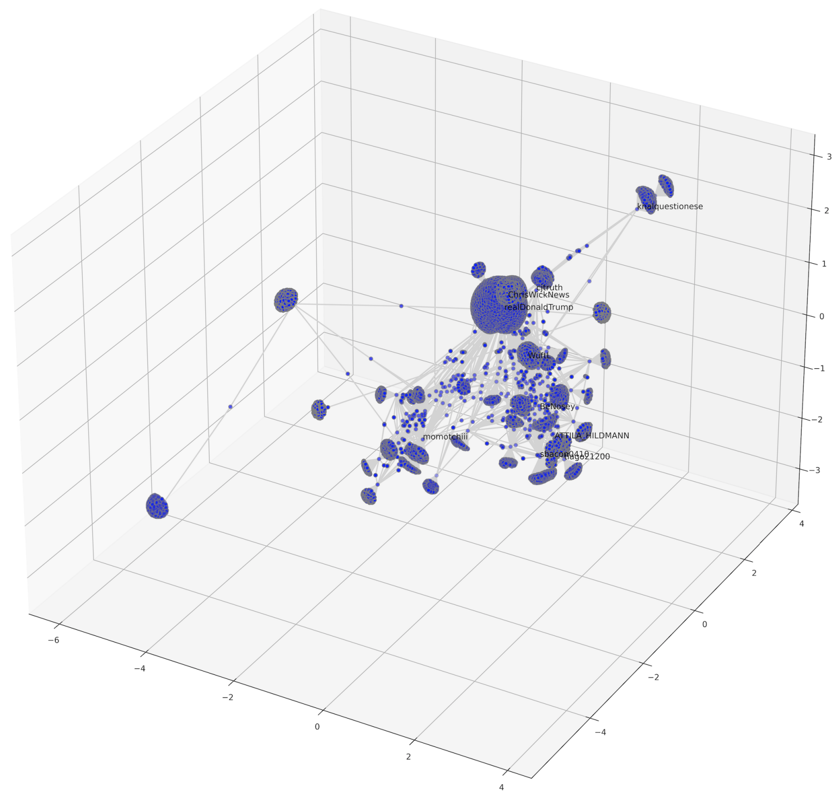

Figure 1: Initial plotting attempts using Matplotlib

3.3 Polarity Graphs

Figure 4 illustrates the entire graph, with nodes coloured based on their polarity measures. Overall it seems like for most nodes we have a neutral stance, either there was no polarity score measured or the tweet is deemed as more or less neutral. This is in line with the polarity distributions mentioned in the earlier section. At a first glance it does however seem that there are more positive (red) nodes in the Trump cluster and several neighbours.

To check this, we can subset the graph into positive nodes only. We then select the largest resulting component and check the main users left in the graph. Figure 5 shows the largest component that is left over when filtering by positive nodes (polarity > 0.01). While this seems to be a much smaller subsection of the bigger picture (Figure 4), it is interesting to observe that the largest cluster remaining is indeed associated with Trump. Since our data is in regards to a specific movement, we assume that positive polarities describe agreement or a general positive attitude towards the *QAnon* movement. Furthermore, since our network seems to be centered around Trump (he has the largest node-degree, as

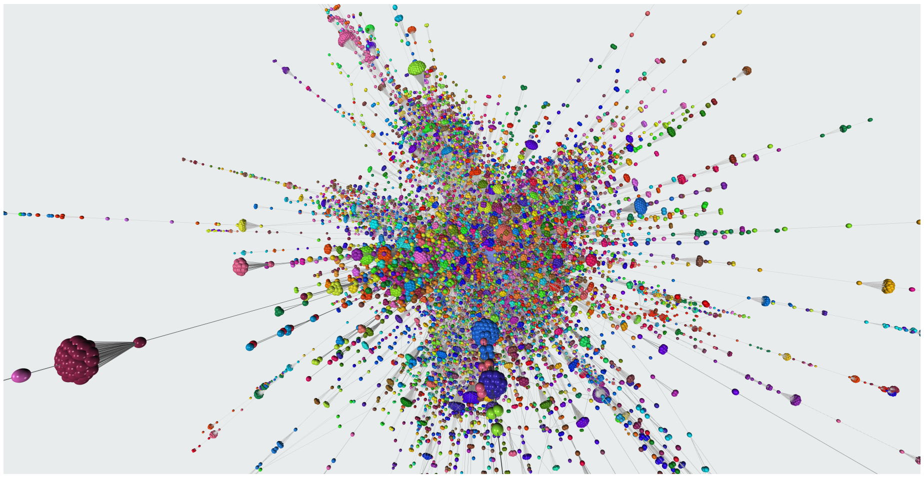

Figure 2: Graph of all users selected based on keywords only (MCL Cluster Colouring)

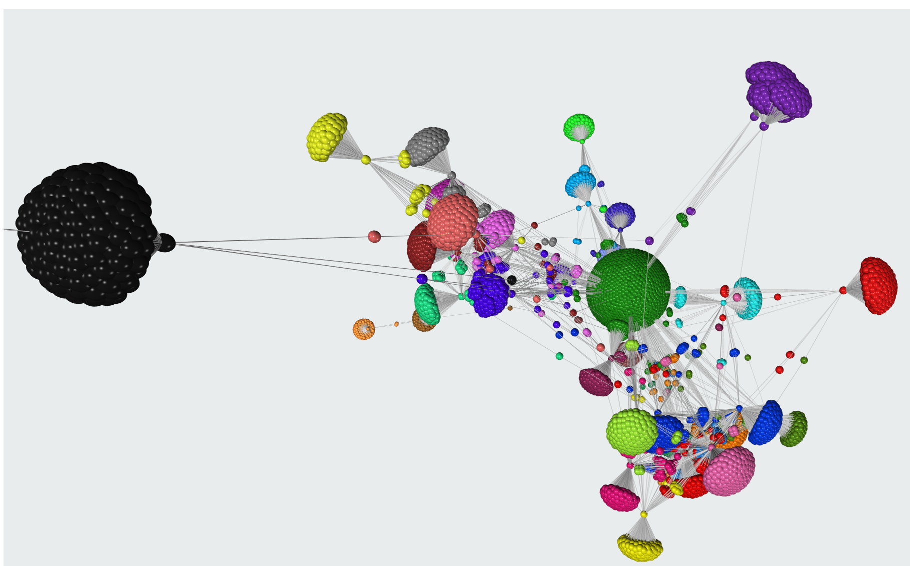

Figure 3: Graph of all users selected based on keywords and top 40 node degree users (MCL Cluster Colouring)

well as largest formed cluster), we also assume that positive polarities indicate an agreement towards the former president. We also see a larger groupings around the **BeNosey**, **SwissBrisq** and **QesWorld** accounts, all of which have now been suspended. The colours of the nodes could also provide an indication of

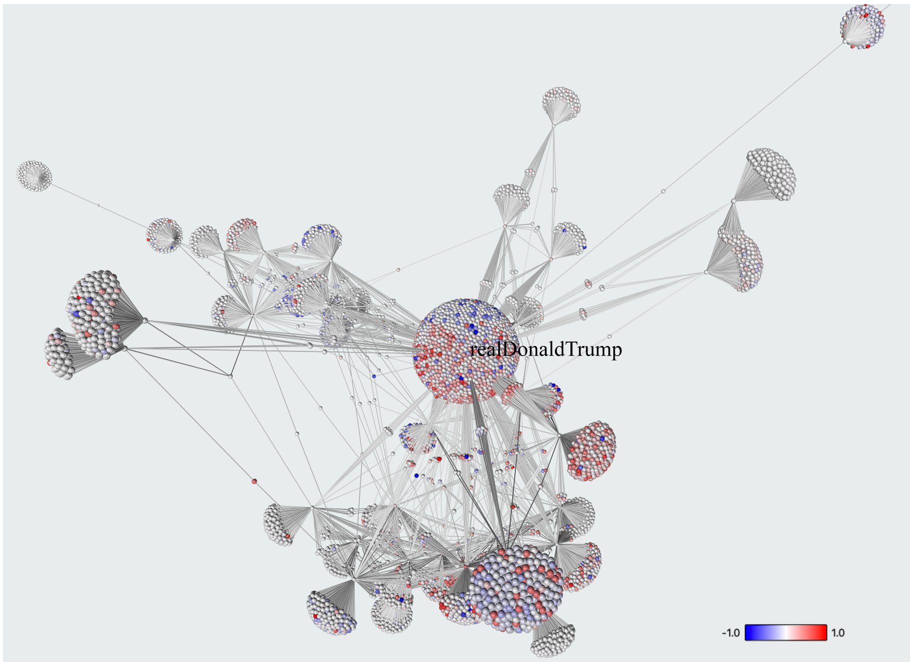

Figure 4: Entire Graph coloured based on polarity scores

the mention relations stances, with more positive nodes (red/orange) observed in the **realDonaldTrump**, **BeNosey** clusters, as well as the connection between **BeNosey** and **QesWorld**, compared to that of **JoeBiden** and **emilyhauser**.

We can also do the same for the negative polarity nodes. As there are less of these in the graph, we do not expect to see very large remaining components. In fact, only keeping negative nodes results in 3644 nodes and 3230 components, meaning most are individual. Nevertheless, Figure 6 illustrates the largest remaining negative components and the main users are highlighted as well.

3.4 Topic Graphs

3.4.1 Full Graph

Using Graphia we can colour the nodes based on the topic to which they belong. This could be useful in revealing clusters of users, and especially source users, who spread information which is of interest. The results of this are displayed in Figure 7. We can already observe some clear dominant colours (topics) in a lot of the clusters. To double check that these findings are meaningful, two users in the "Conspiracy" cluster are selected. Indeed, the relationship is a retweet of the message from the source user **ChrisWickNews** by **WHATCHINGYOUWO1**, with the tweet being *"I DO NOT CONSENT I WILL NOT COMPLY I DO NOT take ORDERS from paid off big pharma connected 'scientists'*

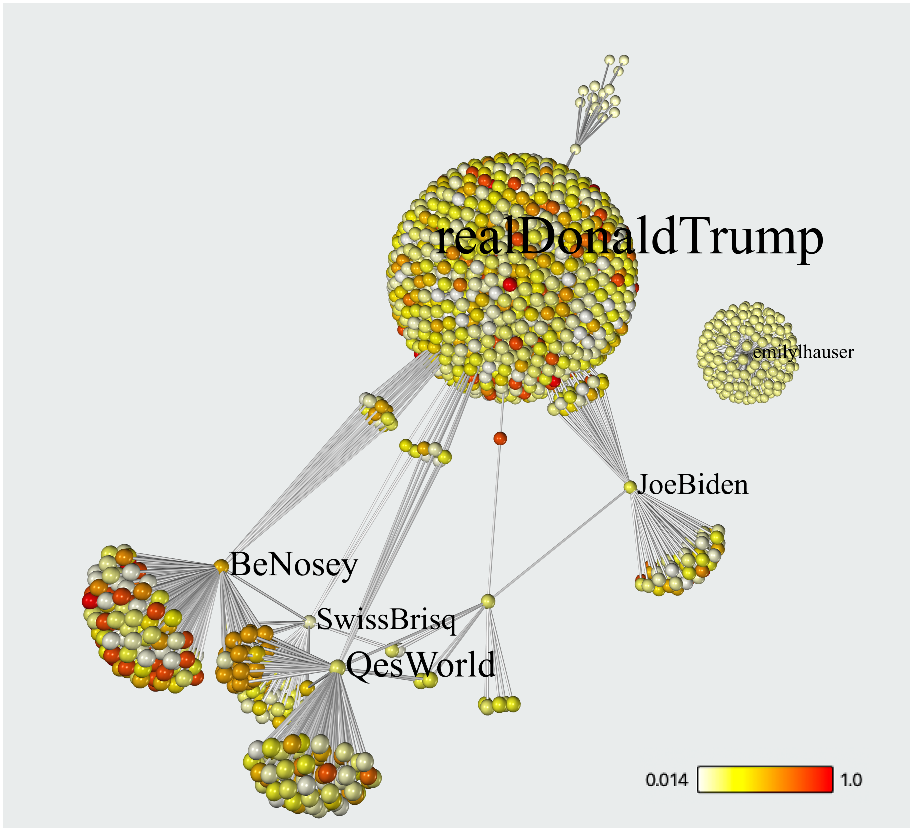

Figure 5: Largest graph component of positive only tweets

with links to billionaire psychopaths. #idonotconsent". The tweet propagates a message concerned with conspiracy beliefs, inline with the assigned topic label.

3.4.2 Contracting Topic Clusters

Another nice feature that we can make use of in Graphia is the ability to collapse/contract nodes which have a common feature. In this case, we can contract by topic and keep all nodes or edges for which the multiplicity measure is larger than 1 (in other words, where some grouping has been formed).

Figure 8 illustrates the result from this process, with the node colour explained by the legend. This is a similar graph to Figure 7, however groups of topics are now represented by a single node, with the size of the node relating to the number of users inside of it. Furthermore, the size of the edge also depends on how many edges the two nodes share, visualising the strength in connection between two topics of conversation. A more thorough interpretation of this graph and the strong topic links observed can be seen in the later discussion

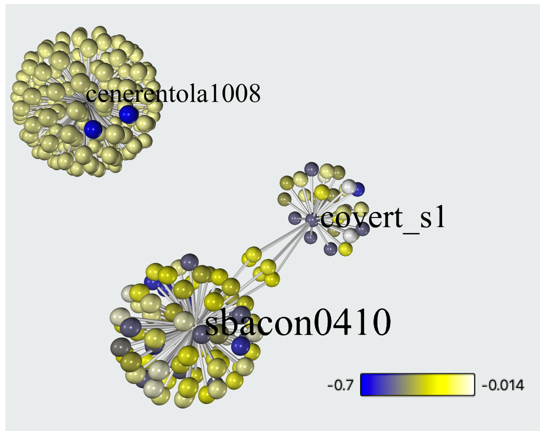

Figure 6: Remaining negative nodes and their largest component

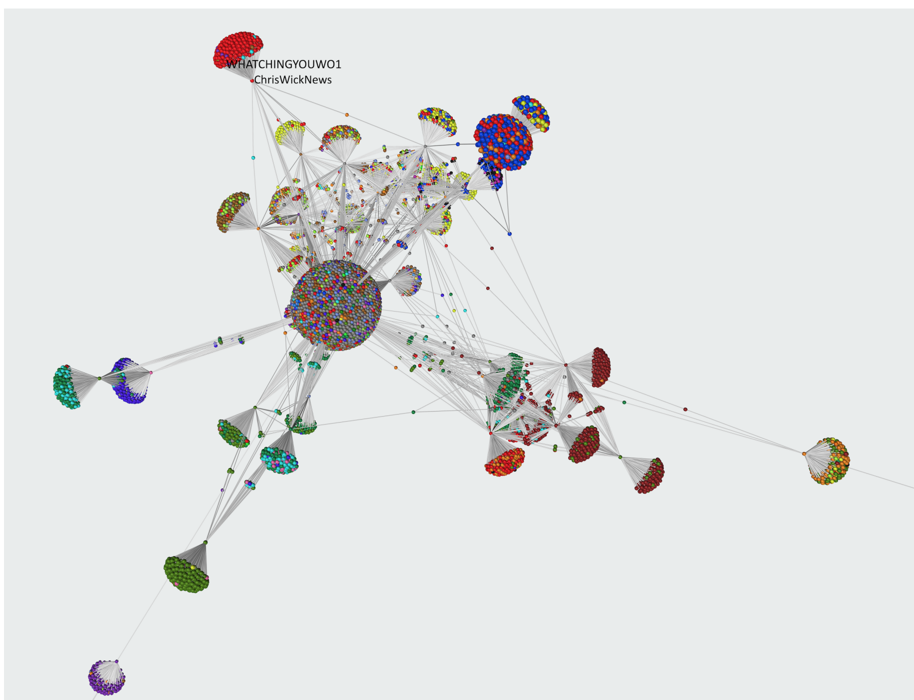

Figure 7: Topic Clusters Visualisation

section.

Selecting individual nodes also reveals the main users within the group, and again highlights the ease with which these popular users can be identified and visualised. This is an important aspect when it comes to identifying the spread

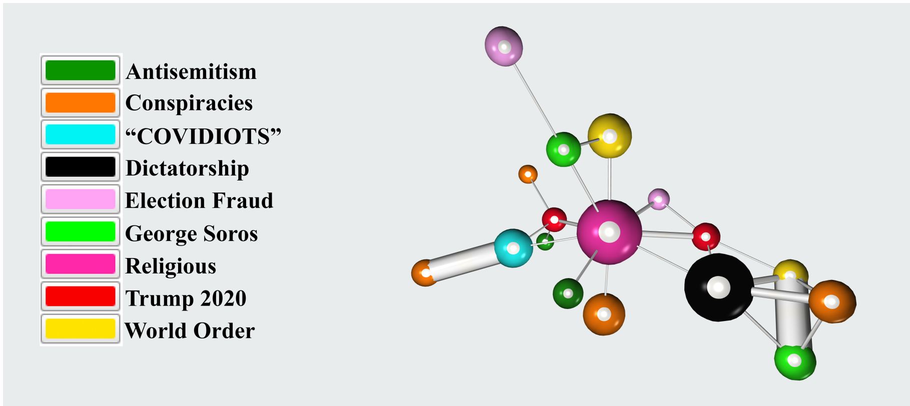

Figure 8: Visualisation of the collapsed topic nodes. Legend on left indicates the topic label for each node.

of "fake-news" and potentially harmful information.

3.5 Location Graphs

Similarly, the nodes in the graph can be collapsed based on the user locations. Instead of giving insight into the flow of topics between users, we can now observe the locations and main areas of spread of the QAnon movement. To do this we first remove all nodes which contain the "Not Found" location. We select the largest resulting component from the subset, and collapse the nodes based on their location attribute. Again, we keep the nodes with a multiplicity of higher than 1. The resulting graph can be observed in Figure 9. The size of the nodes represents the number of users collapsed within the node (we observe that most users are from the USA) and the size of the edge if also representative of how strong that connection is (there are many mode edges between USA and Brazil users compared to that of USA and the Netherlands). Again, we go over the connections and node sizes in more detail in the later sections.

4 Results

Several tables are presented in this section, highlighting in more detail influential users from the clusters discovered in our graphs. Table 3 summarises the main users found in the different topic nodes from the collapsed graph in Figure 8. For each topic, the most influential user is identified. This is based on their node degree (in and out-going connections to other users). For topics such as "Conspiracy" which had multiple separated nodes in the graph, an influential user is identified for each of those. The table also includes a user summary description, as well as the top locations inside the topic clusters. Note that the location "Not Found" is discarded from the percentage calculations.

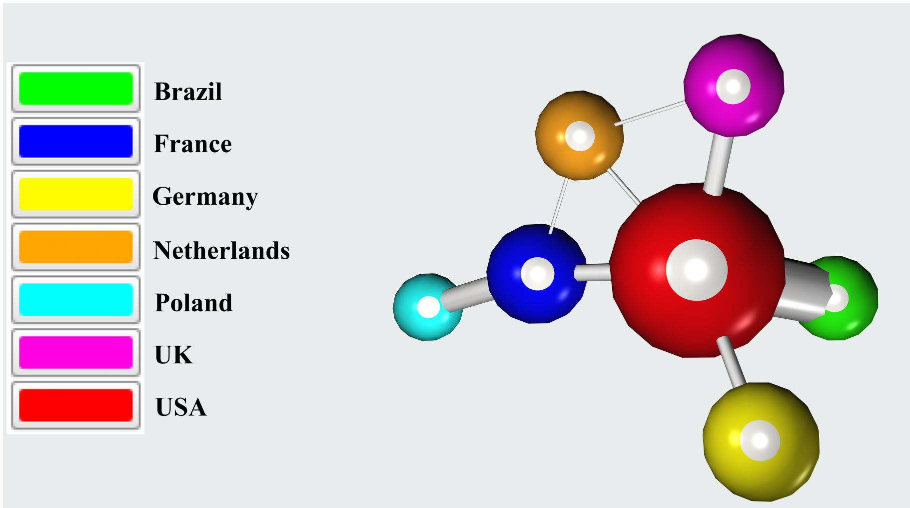

Figure 9: Visualisation of the collapsed location nodes

Table 4 presents something similar, this time however identifying the main users from Figure 9. The table also contains the main topic that is associated with those influential users, and an overall distribution of the main topics at that location.

5 Discussion

In this short project we have demonstrated that the *QAnon* related data collected by Textgain can successfully be used for Social Network Analysis. In particular, using Graphia we were able to create a 3D social network visualisation, as seen in Figure 2. While this graph only included a subset of the original users, the size and "usefulness" of the displayed network does not give meaningful insights into the questions we are aiming to answer. In line with previous work (Borgatti, 2005; Burt et al., 2005) we attempted to make use of graph-based measures, such as the node degree, to reduce the noise in Figure 2. By selecting the top 40 node degree users and the interactions they are involved in, we produce a cleaner graph, Figure 3, where more obvious and easier to follow clusters can be observed (**RQ1**). Nevertheless, using this graph to find users of interest would be tedious and computationally slow work, as one would need to manually scan through and click on nodes they think are potentially influential.

We therefore attempted to include several user and tweet-based attributes to the graph which would be used to group/separate the nodes. Tweet polarity was calculated using Textgain's own AI toolkit[11], the user locations were extracted

[11] https://github.com/textgain/grasp

Table 3: Most influential users based on collapsed topics and associated locations. User accounts in bold have been currently suspended.

Topic	Node Size	Main Users	User Descriptions	Top Locations
Religious	1205	1. **realDonaldTrump**	1. Former US President	USA 74.3% UK 5.4% Australia 2.6% Canada 2.3% France 1.4%
Dictatorship	1087	1. PyrisDeekhlan	1. "Chaotic Writer, AI & Dystopian Worlds". An account with extreme beliefs	France 83.9% USA 2.6% Belgium 2.3% Canada 2.3% Switzerland 1.7%
Conspiracy	1037	1. **AttilaPhotoshop** 2. EChabriere 3. **SwissBrisq** 4. **ForcedAdoption1**	1. Parody account of Attila Hildmann(German far-right conspiracist) 2. French professor in biochemistry 3. Swiss news account dedicated to solving "Deep State Parler" conspiracy 4. Child activist, anti-abortion, religious account	USA 24.9% UK 23.7% France 18.5% Germany 9.7% Canada 3.6%
World Order	757	1. patrick_edery 2. FPatriarcal	1. Editorialist, writer for Polish weekly news outlet 2. "Conservative, Christian, Psychologist."	France 28.3% Spain 23% Chile 17.1% Mexico 3.7% Canada 2.7%
Covidiots	492	1. IngeJusta	1. Conspiracist Dutch account	Germany 56.6% Netherlands 23.5% Belgium 4.1% Denmark 1% USA 1%
George Soros	422	1. MrJumper02 2. **virusdewuhan**	1. Conspiracist French account, against the rich 2. Conspiracist Spanish account, plandemic and Trump	France 64.6% Spain 14.6% Venezuela 2.5% USA 2% Mexico 1.5%
Trump 2020	312	1. JoeBiden 2. suzinator7	1. Current US president 2. Conspiracist US account, anti-mask, pro-Trump	USA 65% France 8.9% UK 4.9% Netherlands 4.9% Australia 2.4%
Antisemitism	55	1. **sbacon0410** 2. HyperionNL1	1. Pro-Trump account and patriot 2. Dutch account, anti racism and discrimination	USA 71.4% Netherlands 14.3% Spain 14.3%
Election Fraud	23	1. **POTUS** 2. omarbula	1. ?? 2. Geopolitical expert, pro-West, patriot, christian	USA 83.3% UK 8.3% Malta 8.3%

from the tweet meta-data and mapped appropriately, and topic modeling was performed on the entire dataset. We were successfully able to include these attributes into our graph and use them to create sub-graphs or find clusters and users of interest (**RQ2**).

Using the tweet polarities we were able to identify and visualise the user interactions in the network which had a positive stance towards the *QAnon* movement. These are all concentrated around the main user in the network, that being Donald Trump, and several other accounts are also highlighted (Figure 5).

Colouring the network based on the topics (Figure 7) reveals a lot of uniformly coloured clusters, without even contracting and grouping nodes based on their attributes. This is a good indication that the method of selecting users based on their node-degree is a strong starting point of attempting to find groups of interest. We went further and contracted the nodes based on their topics, or locations, as seen in Figures 8 and 9. From these graphs we were able to form the results seen in Tables 3 and 4. Both tables provide useful information, and are indeed able to highlight potential users of interest which warrant the need for monitoring. We are able to extract the main users who are propagating the specific topic or tweeting the most at a particular location. Interestingly, certain accounts have already been suspended, likely because of their inappropriate activity. We also observe that the accounts **realDonaldTrump**, **AttilaPho-**

Table 4: Most influential users based on collapsed locations and associated topics. User accounts in bold have been suspended.

Location	Node Size	Main User	User Description	User Topic	Location Topics	
USA	2257	**realDonaldTrump**	Former US President	Religious	Trump 2020 Religious Conspiracy Antisemitism Election Fraud	37.5% 15.4% 10.5% 8.9% 8.1%
France	701	occulture_ytb	Occult Culture YouTubers	Not Found	Dictatorship George Soros Not Found Conspiracy World Order	31.8% 19.5% 18.4% 11.1% 8%
Brazil	263	knalquestionese	Student, Pro Trump and Bolsonaro	COVID Cases	Politics Trump 2020 George Soros Bill Gates World Order	73.4% 9.9% 9.1% 4.2% 2.7%
Germany	250	**AttilaPhotoshop**	Parody account of Attila Hildmann (German far-right conspiracist)	Conspiracy	Covidiots Conspiracy Not Found Football Bill Gates	48.4% 28.4% 12.8% 0.8% 0.8%
Netherlands	85	MarionKoopmans	Dutch virologist and expert	Covidiots	Conspiracy Covidiots Football Trump 2020 Not Found	15.3% 10.6% 8.2% 5.9% 4.7%
UK	49	**suzanne_kerry**	No description	Cryptocurrencies	Conspiracy Trump 2020 News/Announ. Religious Election Fraud	38.8% 12.2% 12.2% 10.2% 4.1%
Poland	8	patrick_edery	Editorialist, writer for Polish weekly news outlet	World Order	George Soros World Order	75% 25%

toshop, **SwissBrisq**, **sbacon0410** and **JoeBiden** come up several times in our different graph/cluster subsets.

Finally, using the contracted graph figures and their resulting tables, we can get an idea of the spread of topics and their location out-reach (**RQ3**). Figure 8 shows the two main topics in the network, that being "Religion" and "Dictatorship". We observe that the rest of the topics and discourse mainly arise from these two starting points. Of the two, "Religion" seems to have stronger links (stronger and thicker edges) to topics such as "Trump 2020", "Antisemitism" and "Covidiots", while "Dictatorship" has particularly strong links to "World Order", "Conspiracies" and "George Soros". We also observe very strong links between "Conspiracies" and "Covidiots", as well as "George Soros" and "World Order". The term "Covidiots" was introduced at the start of the pandemic, to describe individuals who took little notice of the virus spread and warnings, either not believing in its severity or the advice from governmental bodies. Indeed, we expect the observed strong link between them and people talking about conspiracies. George Soros, a famous billionaire, has long been a target of conspiracy theories. The most recent one, highlighting his involvement in the various US protests. As people believe that he, and other powerful individuals, play a part in controlling the public (*New World Order* conspiracy fearing a dictatorship), the findings in our graph are not surprising and support the methods that we applied. This figure has provided a useful visualisation in seeing how the topics develop and which topics (and users subsequently) have strong links and interactions with each other.

A similar discussion can be made about the location spread of the *QAnon* movement seen in Figure 9. As expected, the main location in our dataset is the USA, from where the spread initiates into other countries. Interestingly, the USA has the strongest connections to Brazil, as well as France and the UK. It seems like the spread of *QAnon* has been slightly weaker in the Netherlands, while still being influenced by the neighboring European countries. As mentioned earlier, we observe around 42 source locations in the data, and a much higher number of 142 target locations. This fact, combined with the graph, highlights the speed and spread severity of the *QAnon* movement. The extra bit of information provided by the graph is to indicate where that spread is stronger and thus where future monitoring efforts should be focusing on first.

Finally, we are also able to get some interesting insights from Tables 3 and 4. We observe that certain topics such as "Conspiracy" and "World Order" have a much wider location outreach, with more spread out country distributions. On the other hand, "Religious", "Trump 2020", "Antisemitism" and "Election Fraud" topics are mainly discussed in the US, while "Dictatorship" and "George Soros" are concentrated in France.

6 Conclusion

With the currently presented work, we have shown that Social-Network Analysis (SNA) can successfully be applied to a collection of Twitter data in discovering and visualising clusters of similar topics, locations and influential users. We have also managed to visualise the main discourse flow through the network, as well as the location spread of the *QAnon* movement. We believe the presented findings provide further support for the use of SNA in locating users and groups of people which share a common characteristic and need to potentially be monitored by governing bodies and organisations in the quest to limit the spread of misinformation and harmful fake news.

As mentioned earlier, research into misinformation detection is active, even more so in the current pandemic where false information can be harmful, and the problem is approached from various directions (Guo et al., 2019). Our research presents the basis of an entire pipeline which can be utilized in investigating and detecting misinformation. As the pipeline is reusable and scalable, it can be applied to various contexts. Having general SNA-based pipelines in the future, which can highlight and discover influential users based on some feature, could prove helpful in the need for rapid misinformation detection. Nevertheless, there are various suggestions for future work to build on top of our results.

7 Future Work

For starters, while SeaNMF has been shown to outperform LDA in topic modeling of short texts (Shi et al., 2018), it is still dependent on the number of topics, as well as the pre-processing steps we take (removal of stop words, closed-class words, lemmatization etc). Instead, future attempts could make use of even newer methods, such as Top2Vec (Angelov, 2020)[12]. Unlike SeaNMF, Top2Vec does not require the number of topics to be specified, and is independent of stemming, lemmatization and stop-word removal, while still performing well on short texts. As the two methods approach the problem of modeling differently, it would be interesting to compare the two and see if we are able to locate similar clusters in our dataset. Furthermore, Top2Vec offers more search functionalities and is able to sort topics hierarchically. As of now, our method is able to find several clusters belonging to the same topic, for example "Conspiracy". As this is a general topic, there are likely sub-topics within it relating to different conspiracy types. Top2Vec's search and matching functions, as well as hierarchical structure, could be a way of discerning the sub-topics and getting an even more detailed idea of the discourse distribution in our network.

Another method of comparing what specific clusters are talking about is to apply different text summarisation methods. So instead of first performing topic modeling on the entire data, we locate the clusters based on the influential users we have discovered (high node degrees), and extract all tweets from each cluster. Following that, different text summarisation methods can be applied on the groups of collected tweets. Methods in the field typically relying on comparing word distributions (Villatoro-Tello, Villaseñor-Pineda, & Montes-y Gómez, 2006; Nenkova & McKeown, 2012), with tools such as SAGE[13] having the ability to find the most representative keywords. Successful application could for example reveal differences in the conspiracies that groups of people are discussing, instead of labelling the entire cluster under one general topic.

Another suggestion for future work is to use different graph-based measures for finding and ranking the influential users. Our work simply looked at node-degree, so the number of incoming and outgoing connections. Node-degree is however a local measure, and does not capture the global picture in the graph. Graph-theory introduces a wide selection of measures[14], with each representing a different aspect of the network. For example, we could rank users based on centrality, which takes into consideration neighbors connectivity, or betweenness-based measures which quantify a node's importance in information flow[15]. Another promising measure is Google's PageRank centrality, which introduces a damping factor to control the neighbors' effect while determining a specific node's importance (Gleich, 2015). Attempts to identify users spreading specific topics using PageRank have already been documented (Priyanta, Trisna, & Prayana, 2019), with results leading to different rankings depending on the measure used. Future works could rank influential users based on a combination of the aforementioned measurements which could lead to more valid findings.

[12]https://github.com/ddangelov/Top2Vec
[13]https://github.com/jacobeisenstein/SAGE
[14]http://braph.org/manual/graph-measures/
[15]https://home.kpn.nl/stam7883/graph_introduction.html

References

Angelov, D. (2020). Top2vec: Distributed representations of topics. *arXiv preprint arXiv:2008.09470*.

Bianchi, F., Terragni, S., & Hovy, D. (2020). Pre-training is a hot topic: Contextualized document embeddings improve topic coherence. *arXiv preprint arXiv:2004.03974*.

Borgatti, S. P. (2005). Centrality and network flow. *Social networks*, *27*(1), 55–71.

Burt, R. S., et al. (2005). *Brokerage and closure: An introduction to social capital*. Oxford university press.

Chatfield, A., & Brajawidagda, U. (2012). Twitter tsunami early warning network: a social network analysis of twitter information flows.

Choo, J., Lee, C., Reddy, C. K., & Park, H. (2015). Weakly supervised nonnegative matrix factorization for user-driven clustering. *Data Mining and Knowledge Discovery*, *29*(6), 1598–1621.

Daughton, A. R., & Paul, M. J. (2019). Identifying protective health behaviors on twitter: observational study of travel advisories and zika virus. *Journal of medical Internet research*, *21*(5), e13090.

Ediger, D., Jiang, K., Riedy, J., Bader, D. A., Corley, C., Farber, R., & Reynolds, W. N. (2010). Massive social network analysis: Mining twitter for social good. In *2010 39th international conference on parallel processing* (pp. 583–593).

Gleich, D. F. (2015). Pagerank beyond the web. *siam REVIEW*, *57*(3), 321–363.

Guo, B., Ding, Y., Yao, L., Liang, Y., & Yu, Z. (2019). The future of misinformation detection: new perspectives and trends. *arXiv preprint arXiv:1909.03654*.

Himelboim, I., Smith, M. A., Rainie, L., Shneiderman, B., & Espina, C. (2017). Classifying twitter topic-networks using social network analysis. *Social media+ society*, *3*(1), 2056305117691545.

Hitlin, P., & Olmstead, K. (2018). The Science People see on Social Media. Pew Research Center, March 21, 2018. Retrieved March 21, 2018, from `https://www.pewresearch.org/science/2018/03/21\\/the-science-people-see-on-social-media/`

Househ, M. (2016). Communicating ebola through social media and electronic news media outlets: A cross-sectional study. *Health informatics journal*, *22*(3), 470–478.

Jelodar, H., Wang, Y., Yuan, C., Feng, X., Jiang, X., Li, Y., & Zhao, L. (2019). Latent dirichlet allocation (lda) and topic modeling: models, applications, a survey. *Multimedia Tools and Applications*, *78*(11), 15169–15211.

Kim, H., Choo, J., Kim, J., Reddy, C. K., & Park, H. (2015). Simultaneous discovery of common and discriminative topics via joint nonnegative matrix factorization. In *Proceedings of the 21th acm sigkdd international conference on knowledge discovery and data mining* (pp. 567–576).

LaLone, N., Tapia, A., Zobel, C., Caraega, C., Neppalli, V. K., & Halse, S. (2017). Embracing human noise as resilience indicator: twitter as power grid correlate. *Sustainable and Resilient Infrastructure*, *2*(4), 169–178.

Matsa, K. E., & Shearer, E. (2018). News use across social media platforms 2018— pew research center. *Journalism and Media*.

Nenkova, A., & McKeown, K. (2012). A survey of text summarization techniques. In *Mining text data* (pp. 43–76). Springer.

Ordun, C., Purushotham, S., & Raff, E. (2020). Exploratory analysis of covid-19 tweets using topic modeling, umap, and digraphs. *arXiv preprint arXiv:2005.03082*.

Priyanta, S., Trisna, I. N. P., & Prayana, N. (2019). Social network analysis of twitter to identify issuer of topic using pagerank. *International Journal Of Advanced Computer Science And Applications*, *10*, 107–111.

Shi, T., Kang, K., Choo, J., & Reddy, C. K. (2018). Short-text topic modeling via non-negative matrix factorization enriched with local word-context correlations. In *Proceedings of the 2018 world wide web conference* (pp. 1105–1114).

Villatoro-Tello, E., Villaseñor-Pineda, L., & Montes-y Gómez, M. (2006). Using word sequences for text summarization. In *International conference on text, speech and dialogue* (pp. 293–300).

8 Appendix

8.1 Keywords

A complete list of the keywords used to scrape the tweets for the original full dataset and to subset the number of users in the mentions file.

Keywords: @anjadebruin6, #radicalislam, w w g 1 w g a, #wearetheplan, wwgowga, @jansecorine, jew flu, #sheeple, childeaters, #jewflu, #bilderberg, #deactivatetrump, nazi hippie, @wakkerbelgie, #fakejewsmedia, @der_dritte_weg, #deutschesreich, #onoublierapas, qlobal-change, @tantemarie, sheeple, @driesverhelle, @fashygoy7, #stopthesteal, invasive species, @opendcurrency, occulture, #greatawakening, @reconquistanetz, where we go one we go all, @qw13blchen, @qanonvlaanderen, #nonaq, w—w—g—1—w—g—a, @thebelgianawak1, @kanonne14, #silentwarcontinues, antifa scum, @roeptoetermenno, #q17, #qanuon, @jackbur78928682, @qanonesp, #q, @drsmichakat, #idonotconsent, #awake, #marchfortrump, #aussieq, #jewswillnotreplaceus, @the1111code, #friedensvertrag, @howardm94695166, @grandreveil, #pedogate, @antinwoalliance, #truthseeker, #17anon, @kthe81368618, #plandemic, #thanq, #marchfortrump, #killary, great awakening, #scamdem, #elite, #xrp, w.w.g.1.w.g.a, #questioneverything, #qanonarmy, #digitalsoldiers, #kaiserreich, #deepstate, @anonbelgian, querdenken 711, @nbulondon, @barnabas777, satanic high priestess, #proudboys, #whoisq, #savehumanity , #wwg1wg1, #qanonnederland, @sbacon0410, #qanonuk, @jaap0570, @sunshinequeen1, #守護theplan, bill gates, #wwg1wga, satanic rituals, #antifaterrorists, #kag2020, #sorosfunded, #wwgonewga, #wakeyup, #covidioten, @are_clouds, @n8waechternet, @soja_walton, #millionmagamarch, @r33f3rs, #saveourchildren, #proudboys, deep state, #qarmy, #wearethestorm, #godwins, @rputschke, @realdonaldtrump, #thestorm, #scamdemic, @verbindetpunkte, @nlwake, #qarmy, 4 dimensional chess, end of days, attila hildmann, @qanon_report, #iamq, @germanypatriot, #stopdemachine, @sons_of_victory, @80hoog, @daeno791, #reveillezvous, #qanonpoland, red-pilled, #dictaturesanitaire, @mamavanm, w-w-g-1-w-g-a, pedocrats, @909islive, #darktolight, we are q, @wodanofasgard, #qanonitaly, #wearethenewsnow, @theplandemic, george soros, @anon_decoder, @ntornee, @bloempje141, @bluefirerising, @realdonaldtrump, @covert_s1, @emeralddragons7, @imaugedessturms, @rabbitholewiki, conspirituality, @eldoctormabuse, #thegreatawakeningworldwide, #hcqwerkt, @vlaams_leeuwtje, volkisch, @_alas7or_, #satanic, #queanon, thugs, @robertjensen, #draintheswamp, @realjameswoods, secret government, elders of zion, @wakeyq, @freitweeter, #disneygate, #jfkjr, @elenochle, genius sent by god, @muze211, @freedomwwgiwga, #noforcedvaccines, nostradamus, @wijnand_nl, w-w-g-1-w-g-a, #qnn, @dhs_wolf, #wwg1wgaworldwide, #qanongermany, @freememeskids__, #plannedemic, @benosey, #thegreatawakening, scum, #beprepared, #herecomesthepain, #wakeup, new world order, @truthhitsevery1, @q_far1, holocough, #clintoncrimefamily, ethnonationalism, #reveillonsnous, querfront, @wirsiegen1984, @bingosmurf, satan-worshiping, #maga, @langefrans, blood libel, #qovid, #pedowood, #qanonfrance, bloody purge, @mathijsmennink, #gesara, #proudboys, @wuftl, #agenda2020,

@patriot_empire, #holocough, @dontlookawayor1, #qresearch, @davidicke, @qanonespanol, lutz bachmann, #qanonsworldwide, @marylouq7, #cabal, occult warfare, #hidenbiden, #qresearch , sucharit bhakdi, #epsteindidntkillhimself, #themoreyouknow, #restartleader, armor of god, @davewhitemanwlm, #qanonportugal, @mqspain, #mega, @luetzowq, @menckethurner, @oann, @bunkerrabbit, #savechildrens, sent by god, @songoq3, #wherewegoonewegoall, @yevaava, #asktheq, #qanonspain, #savethechildren, @qesworld, @swissbrisq, #antifaterrorists, #pizzagate, @mrs_leuchtfeuer, @17_wwg1wga, @qaaron369, @tw33tz0r, @baronhertog, #standby, #soros, @hyperionnl1, #agenda21, #thesheeple, @jennerweing, #satanichollywood, @tqm_patriot, #qanon, @whitneyaspers, #stopthesteal, #goedlicht, #youmewe, #stop007, #dontbeafraid, #thinkforyourself, child-eating satanists, maga, covid19, covidioten, soros, proudboys, qanon, agenda21, plandemic, godwins, trump

8.2 Topic Modeling Full Results

Table **??** summarises the complete results of our topic modeling procedure (40 topics, stop-words removed). Each row corresponds to a collection of words representing a discovered topic.

Table 5: Results from topic modeling

Topic Words							Topic Name
covid	mark	virus	microsoft	gane	harvard	vacuna	**Bill Gates**
lonald	president	#trump2020	sent	believe	presidente	hoje	**Trump 2020**
think	sent	truly	thought	thing	days	mind	**Religious**
sind	#corona	auch	#coronavirus	würde	bhakdi	wird	**Covidiots**
ociety	#soros	open	dinheiro	panfleto	brasil	drogas	**George Soros**
andemic	#agenda2030	#covid1984	masks	lockdowns	vaccine	covid	**Conspiracy**
enate	court	fraud	elections	scotus	ballots	investigation	**Election Fraud**
ouvoir	mondial	monde	mondiale	france	ordre	médias	**World Order**
haben	denn	leider	leute	gemacht	sein	schon	**MISC**
nedios	globalista	gobierno	corruptos	quiere	países	través	**World Order**
nichael	querdenken	#berlin	#le0711	berlin	polizei	xavier	**Conspiracy**
dnesday	check	#besafe	week	tips	today	morning	**News/Announcements**
finement2	#giletsjaunes	#confinement	#covid19france	#france	#resistance	#dictature	**Dictatorship**
zion	conspiracy	jews	protocols	qanon	antisemitic	theories	**Antisemitism**
plants	habitat	ecosystem	climate	environment	ecosystems	biodiversity	**Nature**
oresenta	biden	ontem	combate	establishment	corrupto	rico	**Politics**
juste	déjà	genre	moins	dernière	alors	trouve	**MISC**
aharaj	world	india	prophecy	great	leader	leadership	**Conspiracy**
aga2020	#election2020	#kag	#kag2020	#vote	#americafirst	#trumptrain	**Trump 2020**
mpoco	mucho	puedo	puede	seguro	dijo	mejor	**MISC**
pulação	ordem	direita	nova	também	maior	estão	**World Order**
waar	maar	hebben	mensen	heel	goed	zien	**MISC**
uevos	octubre	diciembre	personas	noviembre	medidas	contagios	**COVID cases**
était	part	avait	parents	américaines	faits	jeff	**Amazon**
ripple	#cryptocurrency	#ethereum	#defi	#xlm	#trx	#cryptonews	**Cryptocurrencies**
oessoa	ficar	kkkkk	lembro	muito	porra	gosto	**MISC**
elon	zuckerberg	steve	larry	jeff	bezos	ellison	**Billionaires**
helsea	liverpool	bayern	arsenal	juventus	milan	premier	**Football**
#ltc	#bch	part	swyftx	independentreserve	coinjar	digitalsurge	**Cryptocurrencies**
anni	della	ancora	fatto	dopo	oggi	siamo	**MISC**
#	mini	#	elast	album	photo	concept	**K-pop**
music	radio	tune	#standby	playing	affairs	resist	**Africa**
#tvn	puntos	#canal13	#cafécargado	siendo	superado	falabella	**MISC**
wakeup	excelente	dejo	viernes	gracias	rola	mención	**Radio**
ardzo	dziś	chyba	roku	jestem	czas	maja	**MISC**
kurt	zacarias	cobain	cadela	concluído	balão	phill	**MISC**
blir	mycket	rätt	också	sverige	tror	månader	**Sweden**
pieranie	zorganizowany	lesbos	udział	funduszy	zasilany	turcji	**Organisations**
#rap	#rock	#trap	#jazz	#funk	#poetry	#war	**Music**