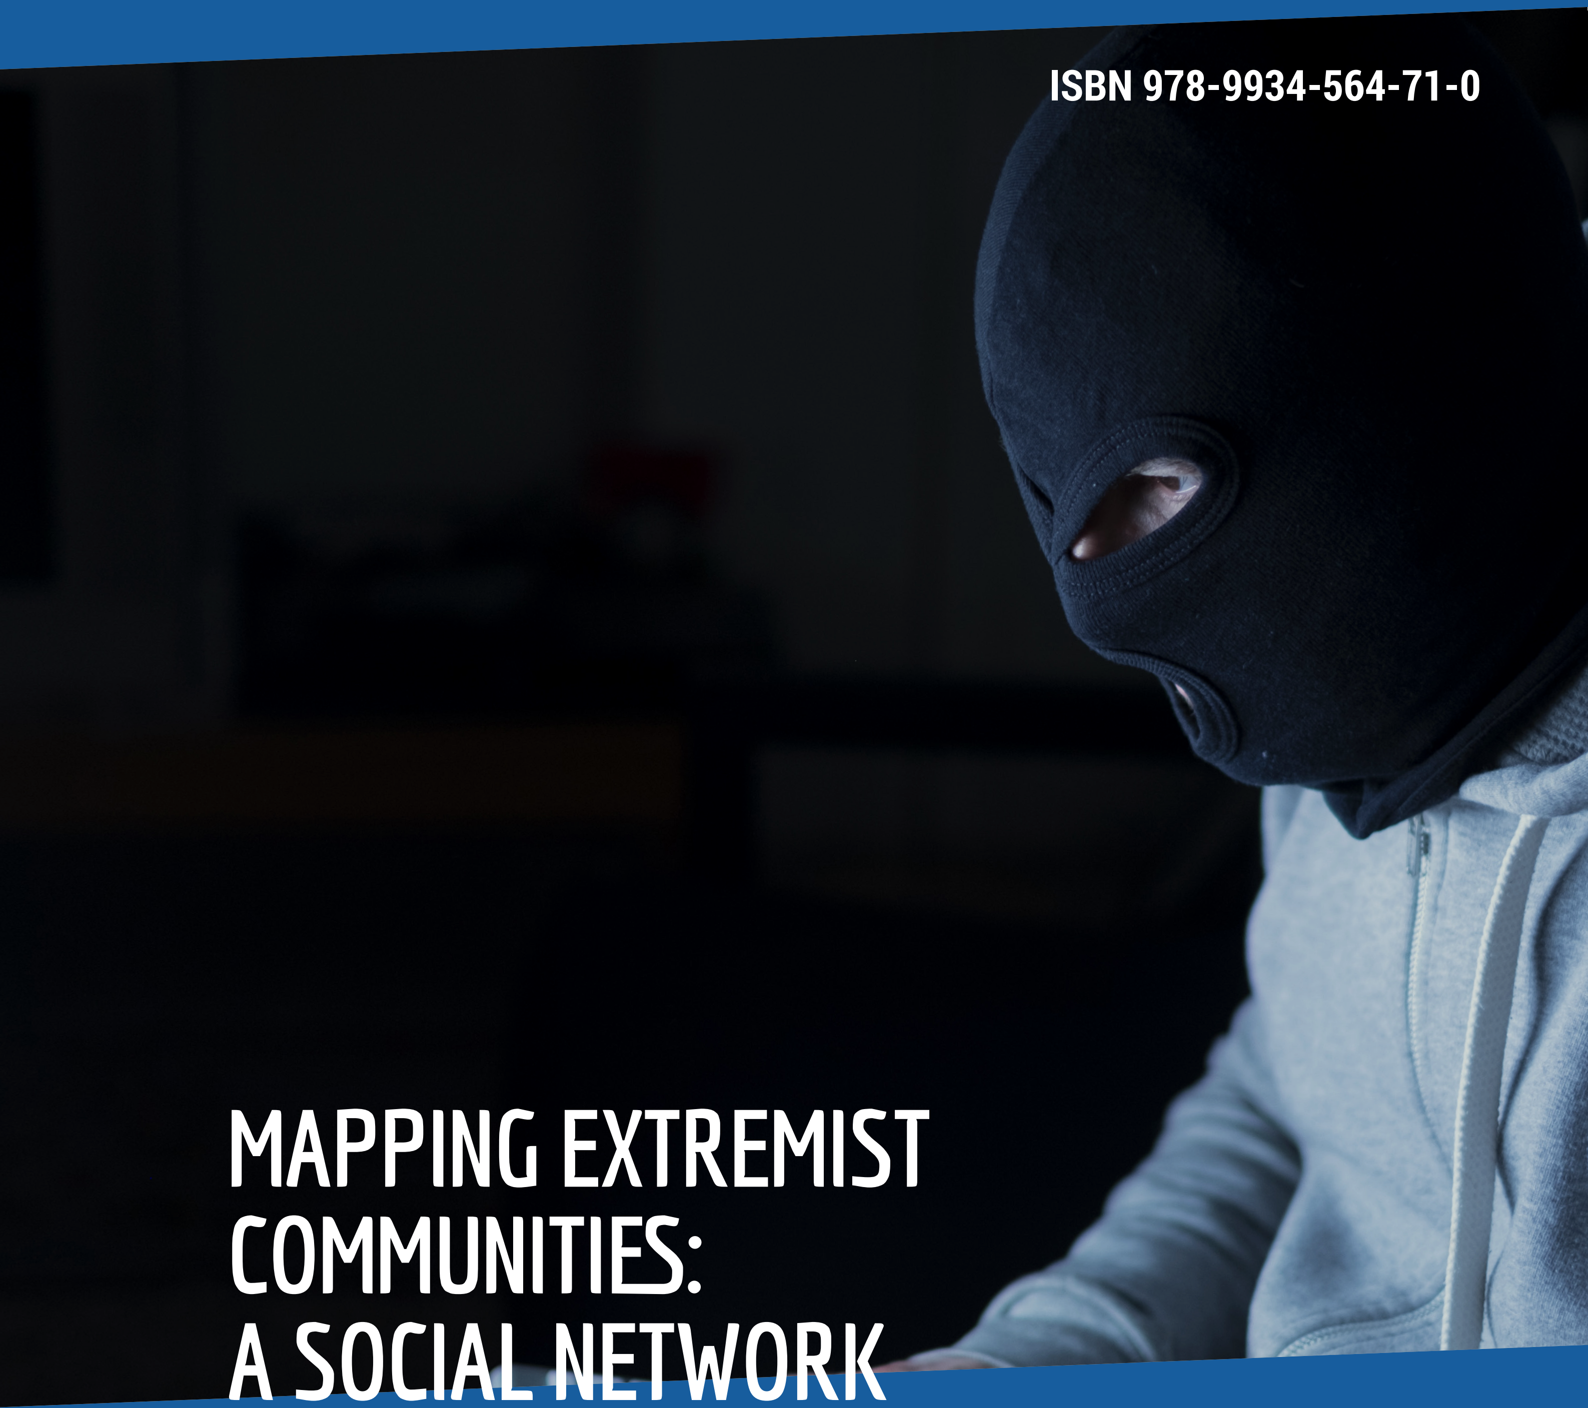

ISBN 978-9934-564-71-0

MAPPING EXTREMIST COMMUNITIES: A SOCIAL NETWORK ANALYSIS APPROACH

Published by the
NATO Strategic Communications
Centre of Excellence

ISBN: 978-9934-564-71-0
Authors: Airbus Defence and Space
Editor: Anna Reynolds
Design: Kārlis Ulmanis

Riga, January 2020
NATO STRATCOM COE
11b Kalciema Iela
Riga LV1048, Latvia
www.stratcomcoe.org
Facebook/stratcomcoe
Twitter: @stratcomcoe

Research by Airbus Defence and Space, socmint@airbus.com

Contents

Extremism is no longer tied to monolithic entities, cohesive groups are no longer the standard.

Executive summary

Online social networks are used by everyone in our everyday lives, including by malicious actors and organisations. Previous work has characterised the specific online behaviour of Middle East-based terror groups.[1] However, this behaviour is constantly evolving, as a response to events such as the battle of Mosul and also due to the strengthening of the platforms' moderation rules.[2] Terror groups target social media platforms such as Twitter, Telegram, and Discord, and while their past behavioural patterns and narrative strategies have been well documented,[3] the adaptive nature of these groups require continuous analysis of their online presence.

A social platform can contain up to two billion accounts (for Facebook), and are a central space where virtual propaganda, recruitment, and discussion happen. During the rise of Daesh, Twitter was used as the online backbone of the organisation's propaganda;[4] more than 100,000 accounts were actively promoting DAESH ideology back in 2014.[5] The combined action of anonymous hackers,[6] improved enforcement of the platform's terms of use,[7] and kinetic military action[8] have greatly reduced this number.

The findings of our study are consistent with those of other research carried out on this topic. In particular, we observe how extremism is no longer tied to monolithic entities, cohesive groups are no longer the standard. We are witnessing a qualitative change—supporters of extremist ideologies are not necessarily active members of an organisation. Extremists individuals do make use of private platforms, but they still are active on mainstream social media platforms. This is a feature which is deeply ingrained in the nature of online propaganda.

In the domain of computer science, the last years have witnessed the improvement of social network analysis at scale. One of the most challenging aspects of social network analysis is community detection; analysts use a variety of tools to visualise the spontaneous group structure emerging from interactions and friendship relations in multi-million-user networks.[9,10] This visualisation, combined with influencer detection and automated text analysis tools such as topic detection, enables the analyst to grasp most of the complexity of a social network.[11]

This computer-science-oriented study explores three lines of research concerning online extremism. First, about the emerging narratives and the topics that can be found on open platforms. We show that many actors actively use terror-group-related terms; most cannot be directly tied to any specific organisation. A second axis concerns the connections between platforms: the information space has no central point as content is shared across platforms. However, the links reveal clusters of locations: we observe a group of Pakistan-India conflict mentions, and a cluster of US alt-right websites, transforming terrorism into a migration problem. The third axis relates to the social media landscape structure. We rely on a combination of document-level topic modelling and graph analysis to detect and explore the social data, visualising the types of groups that are active on the topic. Among the results, we found a small botnet circulating a pro-Daesh pamphlet and a set of grassroot reactions that managed to moderate a controversial pro-Jihadi post on Reddit.

Methodology

Research questions

This article discusses three lines of inquire into online social network analysis:

- **Emerging narratives:** How easy is it to find obvious terrorist messages among today's online social network noise? The first part of our study focused on events that triggered a high usage of jihad-related terms on Twitter.

- **Connections between platforms:** While analysts once had an 'all-Twitter' focus, today we suspect that radical groups hide and share information across various platforms. The second part of our study asks How many external links can we find? How does this new information compare with previous reports?

- **Social media landscape structure:** In the past, online social networks have been used by structured terrorist organisations for propaganda, coordination, and recruitment. Are they still used the same way today? How can we identify the discussions that imply radical accounts? The current trend is for radicals to 'hide', suggesting that open public discussion would be more general, respecting the moderation rules, while the real recruitment and indoctrination would happen in private channels such as on Telegram.

This study poses three questions concerning the collected datasets. The first axis of analysis proposes insights to characterise the current narratives amongst the radicalised-themes, on social media. Then, we investigate the links between the collected social platforms, and the rest of the Web. Finally, we propose to sketch a cartography of social media, through the detection of its constitutive communities, and their characterisation.

Research scope

We limit the scope of our analysis to three datasets, collected from three very different social networks. All data used in this study was publicly accessible at the time of collection; some messages, accounts, or pages have been banned or have since disappeared. We analysed data from three social media platforms – Discord, Reddit and Twitter.

DISCORD

Built mainly for use by video gamers, as players like to chat with each other and coordinate as teams, the Discord[12]chat service enables each user to set up a personal *server*, adding other users by invitation or letting them join freely.

We crawled one public Discord server, obtaining its entire history from April 2018 to January 2019. The discussion space was quite big with 214 users. Although smaller than our tweet corpus, it presents all the features of a true social network: small diameter and skewed degree repartition.

REDDIT

The self-styled *frontpage of the Internet*,[13] Reddit enables its users to post *submissions* with a title, a text, picture, video or URL. To be topically relevant, Reddit is divided into *subreddits*: these are user-created, community-managed pages where submissions are published. Other users can comment on submissions, together these comments resulting in a 'comment tree' that contains all interactions between user accounts for that submission.

We collected the data one *subreddit* from April 2015 to March 2018l it was composed of 942 submissions, which received 5,818 comments from 286 different users.

TWITTER

On the microblogging site Twitter,[14] users post short messages or tweets that are instantly visible to their author's followers, but can also be seen by anyone performing a keyword search.

The tweets analysed in this study were collected through a Stream query requesting data from January 2019 to March 2019 for a set of 14 keywords and 14 users chosen for their relevance with respect to our theme. However, the presence of these words in the user's tweets does not necessarily imply radicalisation. As is often the case, the tweet dataset returned is not guaranteed to be exhaustive, but matches what usual social network analysis tools obtain during their utilisation. Although seemingly limited, our tweet corpus nevertheless consists of more than 400,000 tweets concerning 275,810 accounts (authors and mentioned users).

Terror-related terms	**#Incite_the_Believers**	a propaganda hashtag
	apostate, kufar	designing heretics / deviant thoughts
	aamaq, dabiq, rumiyah	the three main 'press agencies' of DAESH
	caliphate, jihad	obvious references to Daesh narratives
	fard kifayah	'communal obligation', duty to act together as Muslims[15]
General terms related to Islam	**hijr**	'Hijra', the migration of the Prophet; by extension, it also refers to 'migrating to the Caliphate'
	hizb	'Party', 'faction', often used in the name of factions, such as Hizb-ut-Tahrir
	kabba	one of the transliterations of the Kaa'ba, the black stone in Mecca
	khurasan	an old name for the region including parts of Iran, Afghanistan, and Pakistan
Faction-specific terms	**maulana**	the name of a Pakistani terror group leader

To conduct our study, we used keywords already mentioned in the literature. The keywords are presented in this Table, divided in three types: a) terror-related terms, referring to well-known Daesh narratives, b) more general terms related to Islam, included to expand the scope of collection and to investigate the presence of radical groups within this non-radical topic, and c) faction-specific terms. We do not publish user profiles unless they are impersonal accounts—news is OK, people are not.

Detecting Radicalisation Online

Radicalisation detection and characterisation already provoked a great deal of work in various disciplines and contexts. Computer scientists now play a significant role in understanding its online manifestation.

Identifying and analysing online radicalism

The risk of radicalisation through social networks can be investigated by coupling graph algorithms for influence and propagation analysis with text analysis tools that can compute scores indicating radical discourse.[16] This approach has shown promising results when applied to a well-known Kaggle machine learning dataset consisting of 17000 tweets of Daesh sympathisers during 2015.[17]

Using automated techniques to enhance content analysis

A large number of network analysis studies is dedicated to the characterisation of radical online content to help determine whether a new publication expresses radical sentiment and also to better understand the nature of radical messages. The characterisation of radical online content can use automation to expand the scope of the study significantly: as an example, one study collected the 120,000 comments attached to radical, violent videos on YouTube. The study then augmented the information with personal data about the authors of the comments as well as text and sentiment analysis, helping to measure the level of hate and intolerance of the audience.[18]

The analysis of Daesh media reveals specific sentiment features, which are amplified in their key narratives.[19] Textual tools are more difficult to use on videos, but still possible. The main recurrent narratives include demonstration of strength, humiliation of IS enemies (not only the West), the notion of a continuous victory, and the religious righteousness of the group. To obtain keywords and phrases with propaganda elements, one may rely on the following analytical websites:

- clarionproject.org
- jihadology.net
- jihadica.com

These websites aggregate information from various sources (more or less open), without any need to create false identities to gather the radicalised propaganda documents. Media content is often circulated by and discussed among the radicalised communities; new language elements may be detectable from their recurrent analysis.

Tracking terror groups' online footprint

Daesh' mastering of online propaganda shadowed what previously existed; however precedent first-scale threats already used the digital world. The behaviour of al-Qaeda sympathisers was documented in 2010, revealing IP address anonymisation and the manipulation of social media platforms to propagate the *Inspire* magazine. There are obvious similarities between what was documented then and what IS successors are doing now.[20] However, Daesh's competencies reached a higher level in terms of both quality and quantity. Their ability to respond (to moderation, dereferencing, and bans) increased the exigencies for developing more sophisticated counter-terrorism analysis tools.

Contextual 'shadow' of terror-group-related publications

A broader analysis of the use of jihad-related keywords (*caliphate, martyr, crusader, apostate, mujahideen…*) has proven that no one terror group has a monopoly on this theme.[21] Quite the opposite. Such keywords are used by people around the globe with an interest in religious matters: curious people, believers who want to understand the meaning of non-violent *jihad*, as well as extremist but non-violent religious factions. *Hizb ut Tahrir* may be considered such a faction, although a number of countries have listed it as a terror group and banned it.

The takeaway: Specialists already consider terror groups' online, sphere of influence, including media and social media presence, to be as serious a threat as their real-world kinetic operations. Rapidly evolving AI tools are leveraged to enable more rapid and efficient textual, behavioural, and social analysis of terror groups' activities in the virtual domain.

The difficulty lies in assessing the frontier of the adversarial presence, both social and topical: those on the social fringe of a group can become its future (e.g. supporters), or our allies (grassroot opponents).

Social network investigation techniques

To answer questions about online radicalisation it is necessary for researchers to be able to detect interactive topic-focused communities of users on a number of social network platforms, each of which presents site-specific features that are very different from each other. However, every platform produces two types of high-level data—content (here limited to text) and the relationships between users.

Figure 1 below depicts the basic system we use in this study. On the left, is the source of the data—most often content and metadata published on platforms such as Reddit, Discord, or Twitter, consisting of an author's name, the posted message (text), and sometimes that user's relationship with other user accounts. The texts are first processed by a topic detection module enabling us to cluster them into classes or to propose a similarity between them, as shown in the upper part of the figure.

Relationships between users are aggregated into a social graph as shown in the lower part of the figure. This framework allows for the attribution of semantic similarity weights between

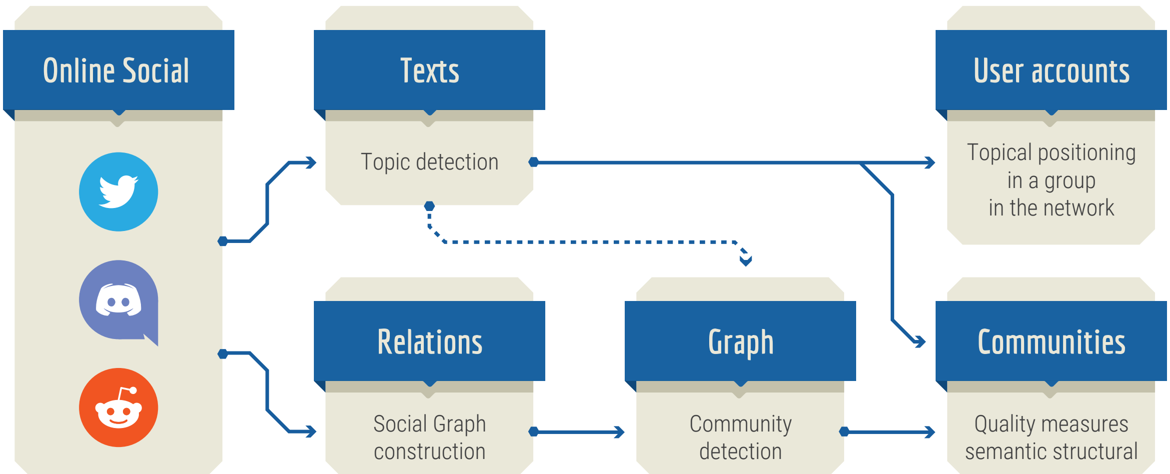

Figure 1. The elements of the social network intelligence system

> Topic detection enables to position users with regard to their groups, or to the network.

users (represented by the dashed arrow); algorithms performing various tasks such as community finding or graph partitioning are used in this step to detect communities of users. Finally, we perform a step of exploitation of the analytics: the topic detection enables to position users with regard to their groups, or to the network; on the community side, we propose measures to grasp their quality, as socio-semantic groups.

Community detection

We are looking for communities based on high levels of interaction between members, which is a relatively well documented problem. The most common approach relies on graph partition algorithms, optimising a community quality value (often the modularity). However, such methods are not always implemented for directed, weighted graphs, and almost never take into account similarities between nodes, to be combined with the presence of edges. We retain only four of them in our experiments: Louvain,[22] Oslom,[23] Link,[24] and LPA.[25,26]

Topic detection and representation

We choses to represent the topics using Doc2Vec embeddings,[27] whose model was trained on a representative corpus of the English language constituted of all Wikipedia articles in English. This solution usually outperforms previous methods such as LSA[28] (latent semantic analysis) in terms of topical precision and scalability.

The difficulty is then to compute a similarity between the messages emitted by two different users. Let u_a be the set of document vectors $d_{a,i}$ published by user *a*. The common similarity measure between two documents is the cosine similarity, denoted cos_{sim}. To compare two sets of vectors, we use the similarity of their average positions, which makes sense in the retained semantic vectorial space and dramatically reduces computation times.

Emerging narratives in a radicalised context

Jihadi terror groups are not only Syria-based

Our first insight from studying the Twitter dataset came from a peak in activity on 3 March 2019—5000 messages with jihad-related keywords commenting on the rumour about the death of Masood Azhar, founder of the Jaish-e-Mohammed terror group in Pakistan, arose in a single hour.

Examples of such tweets include:

- *RT @TimesNow: #BREAKING Reports suggest that Maulana Masood Azhar is dead. | Reports yet to be confirmed.*
  quite objective, this retweet of a post from *TimesNow* ['India's most-watched English news channel'] signals uncertainty about this death.
- *'India has ensured we start our Jihad against it,' brother of Masood Azhar confirms Indian strikes against [Jaish-e-Mohammed].*
  the only pro-jihad declaration is attributed to the brother of the founder, and directly targets India.
- *"Masood Azhar by the grace of Allah, is safe, sound and alright" - JeM statement That's what we wanted to confirm.*
  Denying the information, according to JeM - the terror group in question.

This peak in activity reminded us that Jihadi terror groups are present in a much broader area than Syria and Iraq.[29]

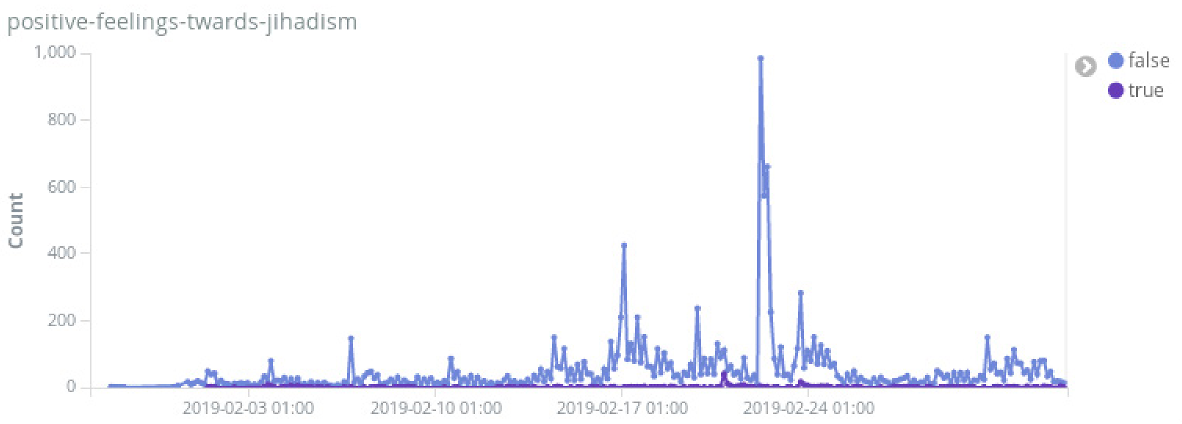

Figure 2: quantity of tweets with/without positive sentiment towards 'caliphate'

> The war against terrorist organisations is also the theatre of Western politics.

Not only jihadists talk about the caliphate

The 'last battle' against DAESH took place at the time of data collection: remaining ISIL forces were defending Baghuz and the coalition was concentrating its efforts there.[30] At this time the Western media were writing about the end of the caliphate as a territorial entity. The use of the term 'caliphate' was massive at this moment, though mostly by informative sources.

This 'final victory' occurred on 22 February 2019, the highest peak in the time series graph above. The represented signal is the number of tweets containing the keyword 'caliphate'; tweets exhibiting positive sentiment towards jihadism are represented in dark blue, while the tweets represented by the light blue line do not exhibit positive sentiment—these tweets are either informative and neutral or centred on the war and its victims, exhibiting negative sentiment.

Our investigation of the URLs exchanged in the datasets collected from the three social networks confirms that the war against terrorist organisations is also the theatre of Western politics. One of the most visible stories at this time, 'DAESH Bride Vows Son "Will Grow Up to Be a Jihadist" as Crumbling Caliphate Evacuated',[31] was promoted by the far-right news website Breitbart.

Interconnections in the social web

The social web is built around content and link sharing: the volume and diversity of the referenced websites is of utmost interest. This section reviews the most significant insights learnt from analysing the URLs shared across our three datasets.

Beginning with the **Twitter dataset**, Figure 3. lists the most commonly shared website domains (i.e. aggregating various pages). Here these are mainly press and news articles—leading global sources such as *The Guardian*, *CNN,* and *Reuters*, but also local press titles (*La Voix du Nord*, *Times of India*). The most cited domain is Twitter itself: users retweet messages or posts links to other accounts. Other social websites and video sharing platforms such as *DailyMotion*, *YouTube*, *Periscope* are referenced, as are several link shorteners (bit.ly, woolsay.com, trib.al).

More surprisingly, engaged 'alt-right' political websites appear in the list: *Breitbart*,

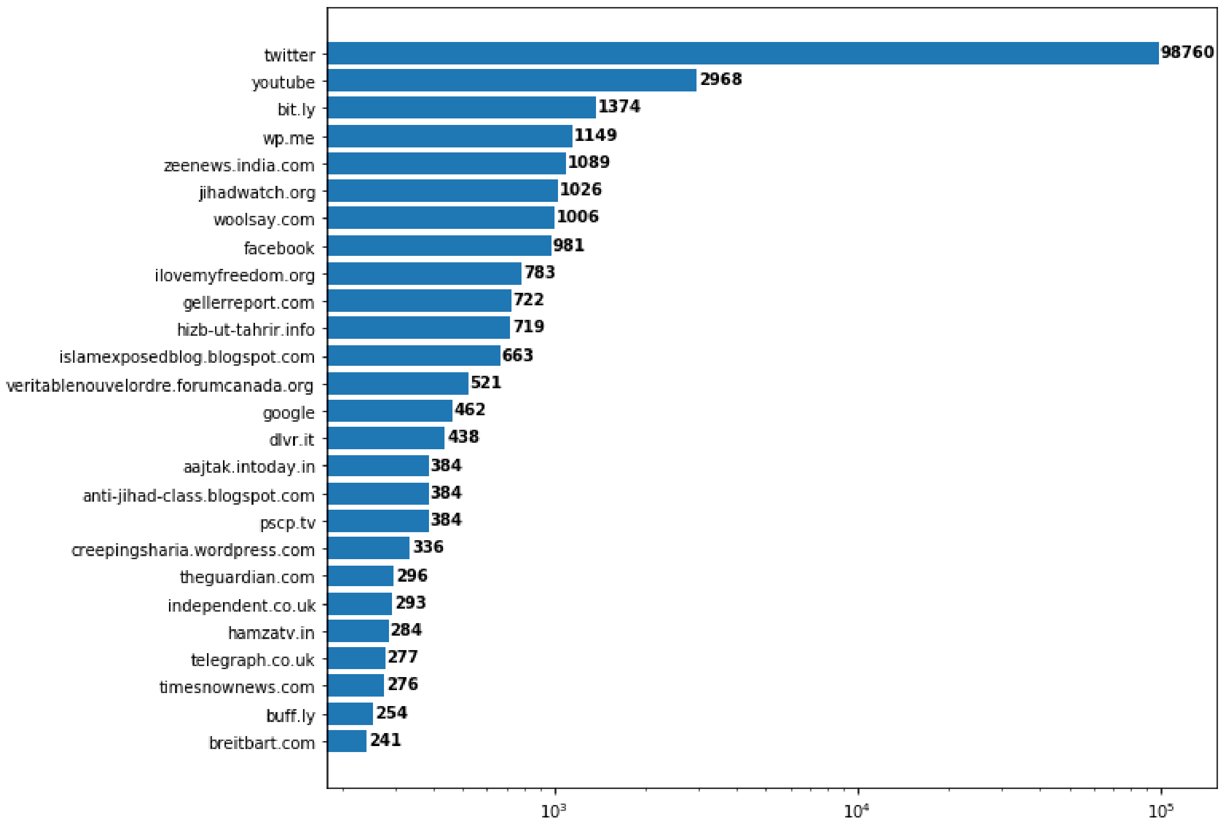

Figure 3. Most frequent web domains in the Twitter dataset (log scale)

jihadilhan as an alias for *Culttture*, and the *Gatestone Institute* are first-rank sources of jihad-related articles.

The links in the **Reddit corpus** are quite different as the collection was not performed using keywords; instead we selected a subreddit containing Salafi contents and views. Sites as *asharis* or *authenticTauheed* publicly remind 'believers' to strictly respect the Content Policy the rules; they do not seem to recommend joining a terrorist organisation.

In the **Discord dataset**, users often link to *Discord.gg* itself, either communicating on the same server or contacting others. The two other significant domains are *YouTube* and *SoundCloud*, both of which host media content—video and audio files. In the case of *SoundCloud*, the links are mostly distributed in a channel called *nasheed* (a vocal music with elements of prayer).

Twitter offers another way of sharing links—instead of including them in a tweet, users can write a link into the **'description'** field of the profile itself. This field accepts text, and it is visible to anyone viewing that user's profile.

Along with the commonly shared links (e.g. Facebook, Twitter, Google, Amazon,

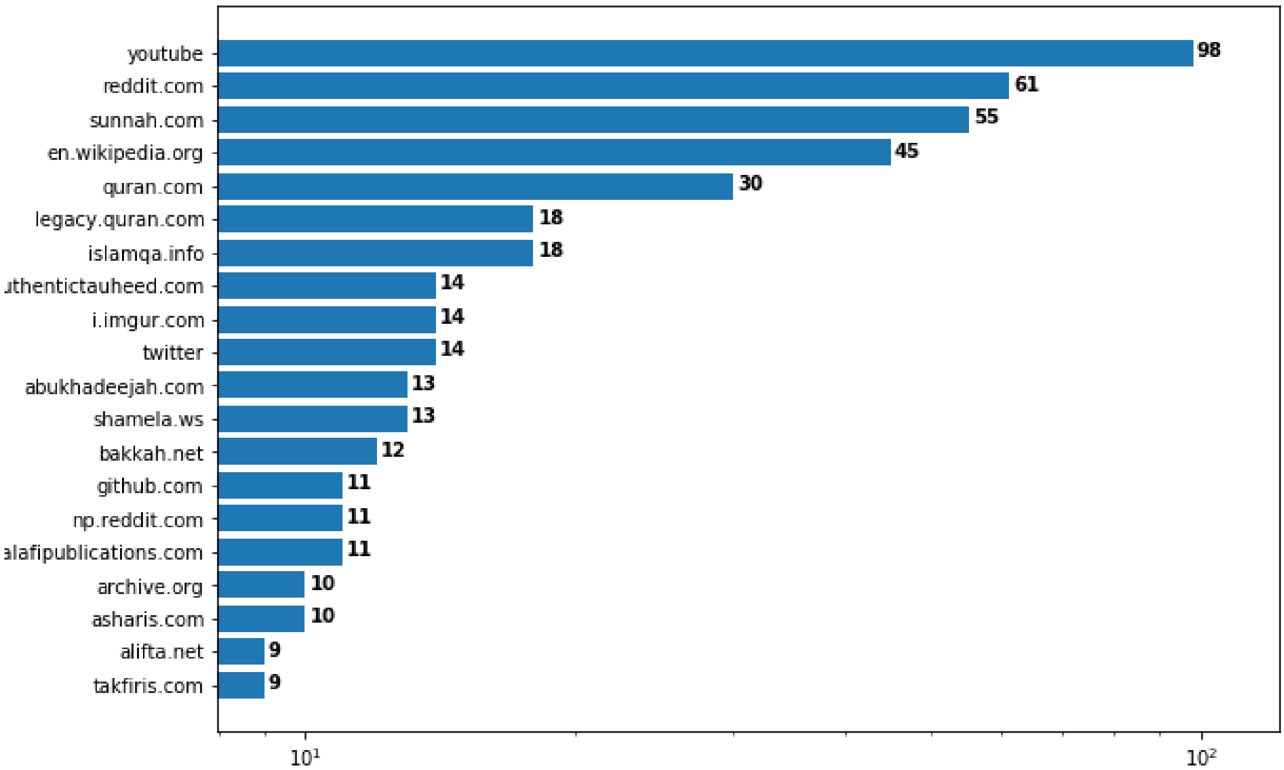

Figure 4. Most frequent outgoing links in the Reddit dataset

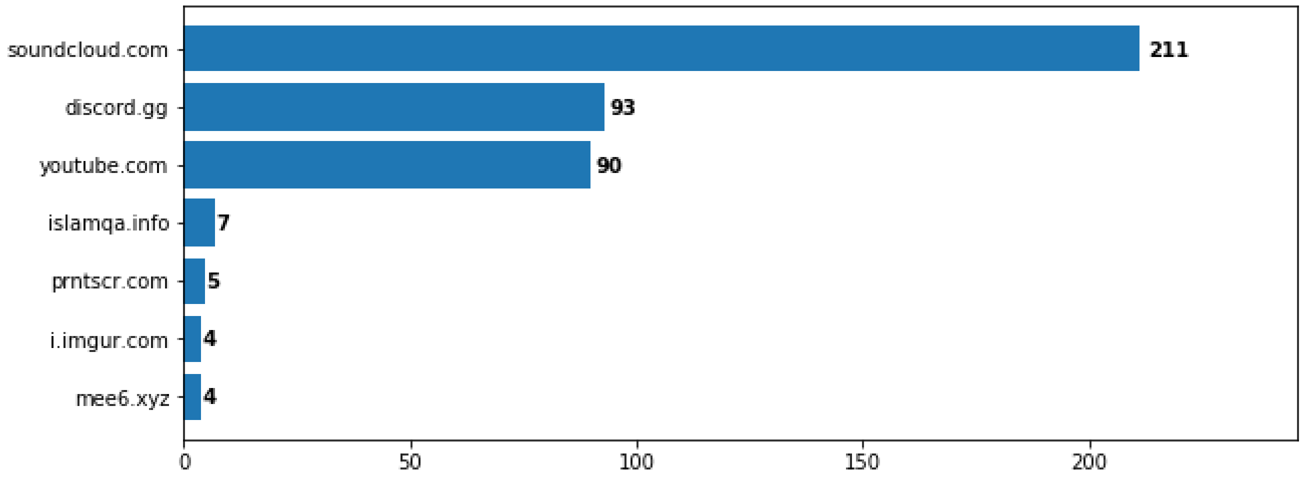

Figure 5. Outgoing links in the Discord dataset

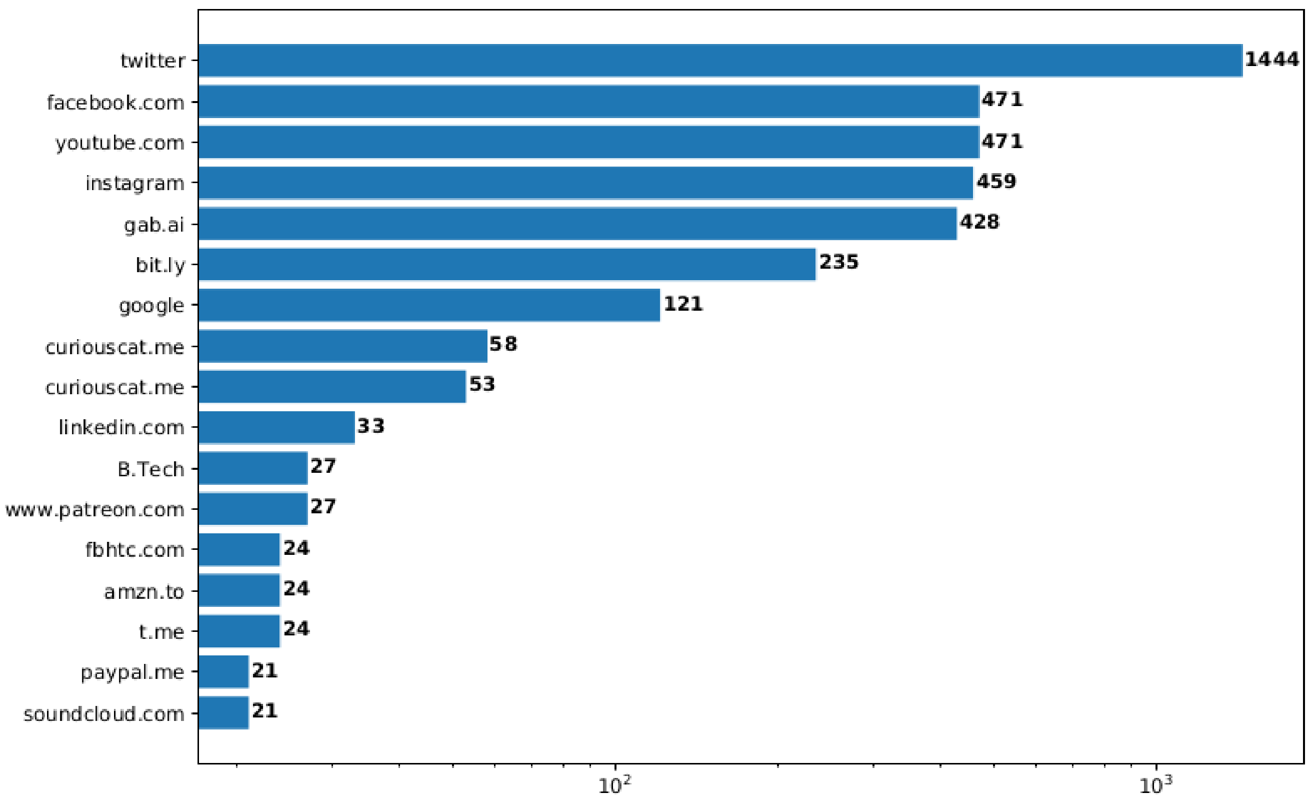

Figure 6. Links contained in the personal 'description' fields of Twitter users

> The 'group' structure is often credited with exerting a strong social influence, even without having strong leaders.

Patreon, etc.), and links going back to Twitter, in this dataset we also find *gab.ai*. The latter is a microblogging platform that 'champions free speech', but which is often seen as a space for alt-right accounts banned from Twitter, and *curiouscat.me*, which enables people to ask questions anonymously. Such anonymity makes it much safer for recruiters to make first contact with potential recruits for extremist groups.

Previous work done by the NATO Strategic Communications Centre of Excellence confirmed the central role of Twitter as an 'umbrella platform that connects the various sources into one easily searchable, browsable information index',[32] mostly aggregating links from YouTube, JustPaste.it, and the Internet Archive digital library. In our case, data collection was based only on keywords and so did not guarantee the allegiance of the users.

However, **interconnections between social platforms have greatly increased** (links to Facebook, Instagram, and other smaller platforms), suggesting that there is no single central point for content sharing. Moreover, file hosting platforms are no longer at the top of the list: first, because they are masked by link shorteners (either bit.ly or Twitter's internal service t.co), and second because obviously radical content is mostly shared on private platforms. Instead, news websites and forums are at the top of the list.

Communities

We visualised the social network through its apparent communities: the 'group' structure is often credited with exerting a strong social influence,[33] even without having strong leaders. We used our datasets to build a network of interactions between the users. On this network/graph are then applied 'community detection algorithms', which find strongly connected sets of nodes/users: the communities, also called groups in the following.

As tweets are far more popular than the other social media, the resulting dataset is bigger than the previous ones. The graph construction is made in two steps. To build the graph from the messages, an edge (u,v) is added if the author *u* of a tweet mentions a user *v* in the text: as such, it considers equally the replies, retweets, quotes and simple mentions. In a second step, we removed non-reciprocal edges to limit the impact of ``random" one-time contact, which is not to be considered an interaction between two user accounts. Doing so, we also removed isolated nodes, as they are not relevant for community detection: the final graph contains 6,276 nodes and 9,721 edges.

Community detection and characterisation

Community quality scores are shown in Figure 7. The horizontal axis, 'Semantic internal similarity', measures the topical dispersion of the members of a group. A low value implies that the users attached to the community produce text messages about a large variety of topics; a high value means the texts are topic-focused and employ terms with similar meanings. The vertical axis, TPR, measures the proportion of triangles inside the group. Intuitively, well-integrated groups contain a high number of nodes in relationships with other group members, while hierarchical and isolated structures contain only a low number of such triangles.

Figure 7 compares four different algorithms used to detect communities, each of which has different strengths and weaknesses. The *link* algorithm detects smaller, more cohesive groups; *louvain* has a tendency to find very large groups and is best suited for large-scale information propagation studies.

Zooming in on three different groups

We now take a closer look at three Twitter communities detected by the *LPA* method to help visualise the meaning of the quality measures. The groups A, B, and C are the yellow points on the lower-left, lower-right, and upper-right corners in Figure 7, respectively.

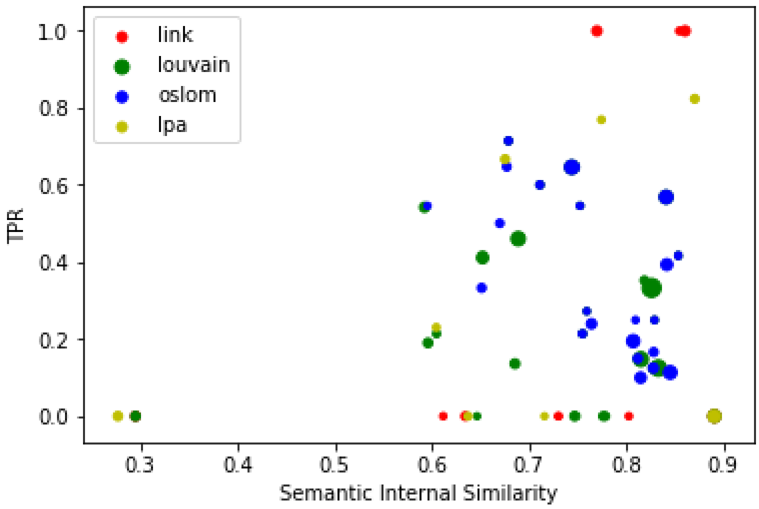

Figure 7. Semantic similarity *vs* triangle ratio of detected interactive communities on Twitter

The groups themselves are represented in Figure 8. Groups A and B are star-shaped, resulting in a zero TPR score, while group C shows much more exchange between its members, much farther of a 'retweet' group scheme.

The yellow group is mainly focused on the dissemination of an essay entitled *Is the Caliphate Irrelevant in the Modern World?*; most of its members have been banned since data collection. It matches a high semantic similarity (most accounts being relays of the publication of the essay) with a low TPR, as the accounts propagate the link to their followers without mentioning other accounts.

This group seems refers to 'Hizb ut Tahrir Indonesia' as it argues about the unicity of the Caliphate and the place of Indonesian Muslims inside it.

Topical view of communities

In this section, we propose to represent the topics discussed by social network users, in a 2D illustration. Even though there are many ways to propose measuring the distance between topics, our choice of mathematical tools (first, Doc2Vec to have a vectorial representation of documents; second, a spatialisation method) provide a well-grounded approach.

The *t-SNE* method[34] produces new axis (here x_1 and x_2), which do not bear semantics but are deemed to best represent the distance between data points. Blue dots correspond to 2D projections of text representations (doc2vec embeddings) for a random subset of 2000 users of the dataset, enabling to illustrate the context in a broad scale.

At a first glance the three communities appear relatively grouped, each one with its colour. However, a closer look at Figure 9 reveals the loose grouping of group A [red]; even though they have outliers, groups B [yellow] and C [black] are much more consolidated.

Group C seems to belong to a blue island; it is related to the language used in their tweets, written in Malay. Various factors draw our attention towards this group: the frequency of mention of the term 'jihad', the strong topical cohesion of the community and its very structure: two nodes are at the centre, without any single-point of failure. The group indeed debates and cherishes 'Ustaz Jihad', a secondary character of a Malaysian soap opera...

On Reddit, a similar process

This study does not pretend to be exhaustive; however, a similar process, applied to the Reddit corpus, let us document new findings about the virtual discussion around Islamism. The Oslom algorithm spotted a community with high semantic cohesiveness (0.67), and low conductance (i.e. the group does not speak with the rest of the network), which indicates a one-shot discussion. The group is visible in Figure

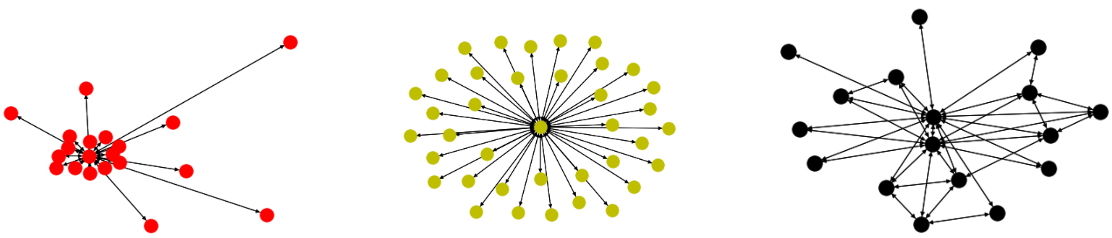

Figure 8: Groups A, B, and C:—three structurally different communities identified by their tweet interactions

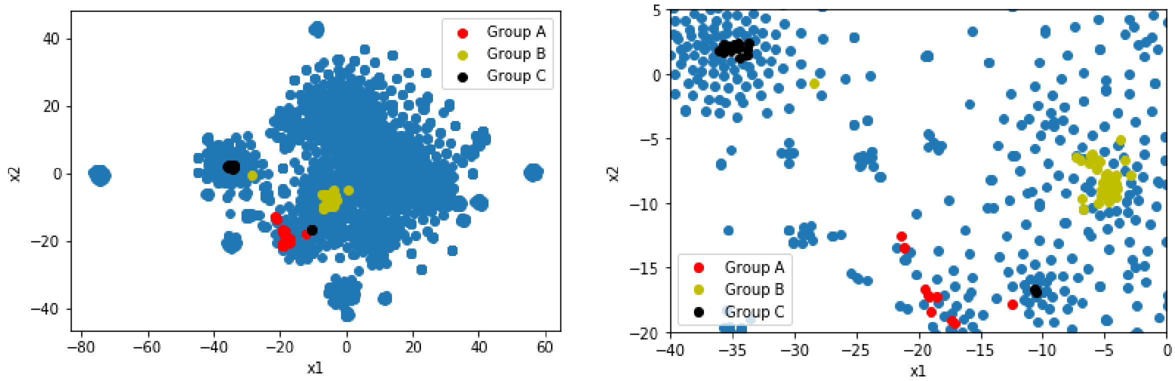

Figures 9 a) and b): on Twitter, topical positioning of users from the 3 colored groups; zoomed on the right

> Messages ranging from polite disagreement to more engaged and provocative statements seeking to lure the other into a debate.

10, where nodes are coloured according to our interpretation of their stance. The green node triggered a debate by sharing links from *IslamQA*, notably around the death penalty for ex-Muslims who are condemned for renouncing their faith.

The blue nodes are referred to as 'innovators' by most of the users of this subreddit: they appear to give some credit to Western values such as human rights, state justice, and scientific inquiry. Green's stance is opposed by messages ranging from polite disagreement [*Hey, I don't think that is really nice to post.*] to more engaged and provocative statements seeking to lure the other into a debate [*Why didn't you respond to his point about a non-Muslim country executing someone for converting to Islam?*].

These (supposed) grassroots respondents appear in groups when an initial message contradicts human rights, but are absent during more abstract debates—they do not seem to be trolls or Islamophobes.

Conclusion of the social network analysis methods

Beyond the limited number of examples presented in this paper, the community characterisation algorithm aims to build a situation map of the analysed online social network. Instead of an unending flow of messages, the expert now dives into a smaller number of social groups, for which social and topical cohesion are measured; moreover, these groups have been located topically, which may accelerate our evolving social media cartography.

The power of these tools should not lead us to forget the basics: it is difficult to eliminate bias at any stage of a study, from data collection to model construction. In our case, the social links between individuals are surmised from the explicit activity logged in the network and the semantic contents of messages are interpreted through a Wikipedia model, either of which may prove inadequate in some specific cases.

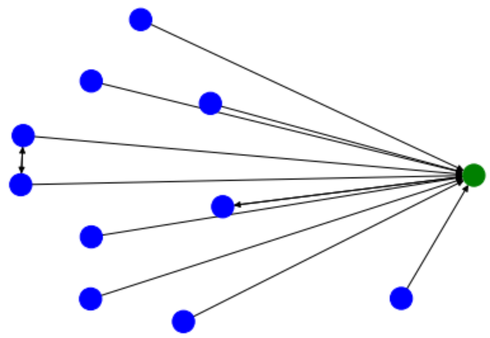

Figure 10: A community hosts a controversy, one against all

A model of the 'radical' media space structure

Literature inspired us to come up with this schematic figure as a working hypothesis about the way radical groups use the information space. It considers Twitter as the central exchange place, where clean accounts promote the philosophical and acceptable ideological points of a Jihadi ideology, while also proposing entry points to the hidden social web through Telegram channels and other discussion spaces. Propaganda documents and videos (e.g. *Aamaq*, *Dabiq*) that are hosted and republished on a variety of websites, are still promoted and relayed by short-lived accounts on Twitter at the price of being continuously banned.

Our analysis indicated that that Twitter does not obviously play a central role in radicalisation and recruitment. While it still is a space where the activities of jihadi and terror groups are actively discussed, it does not provide a 'discussion space' for the terror groups themselves. Instead, the radicalised discourse takes place on other platforms where moderation is not as effective.

Moreover, our research shows that jihad-related terms are used by traditional media when reporting on emerging events, by fake news outlets exploiting this emotion-laden topic to promote their views, by Western alt-right factions reinforcing the perception of Islam and Islamism as a first-order problem, and by grassroots activists confronting extremist views through open debate.

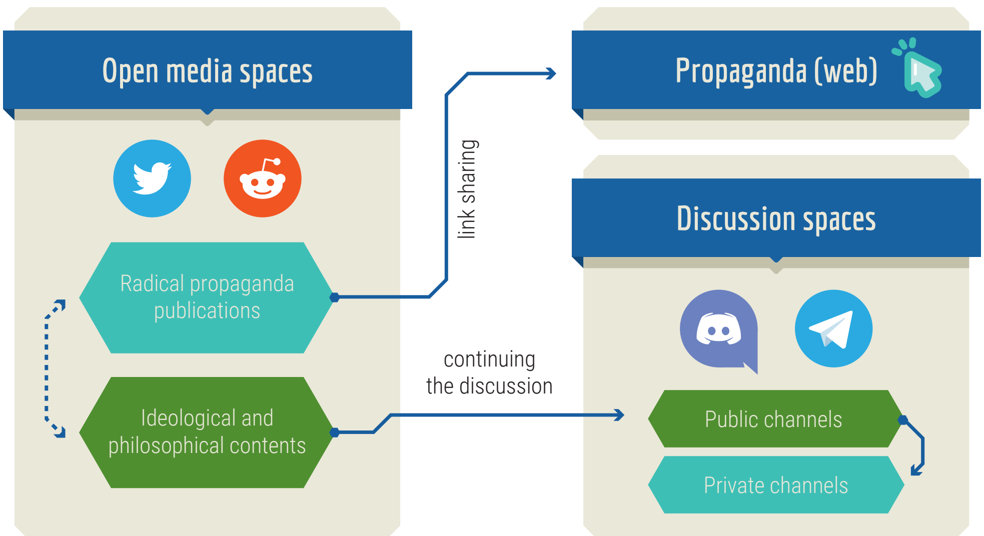

Figure 11. Hypothetical structure of the 'radical' media space

Conclusions

We observe that extremism is no longer tied to monolithic entities (parties, movements, organizations), cohesive groups are no longer the standard. We are witnessing a qualitative change—supporters of extremist ideologies are not necessarily active members of an organisation. Extremists individuals do make use of private platforms, but they still are active on mainstream social media platforms.

Many actors actively use terror-group-related terms; most cannot be directly tied to any specific organisation. However, the links reveal clusters of locations: we observe a group of Pakistan-India conflict mentions, and a cluster of US alt-right websites, transforming terrorism into a migration problem. We also found a small botnet circulating a pro-Daesh pamphlet and a set of grassroot reactions that effectively moderated a controversial pro-jihadi post on Reddit.

The complexity of online social media constitutes a challenge: the presence of malicious entities is difficult to detect and to qualify. In this article, we focused on terror groups based in the Middle East and their main narratives: it is obviously important to update our perception of their online presence.

During the case study, we underlined the presence of an open 'philosophical' discussion that may be an entry point to less acceptable private channels, hosted on non-public platforms such as Telegram or Discord. We identified a strong alt-right and fake news website presence in the themes we investigated, which boosts the Jihadi-related footprint in the public debate, perhaps with the goal of triggering a reaction from their supporters. Finally, on Reddit, an open media platform with a strong comment structure, opponents to Islamism appear with their views and arguments and seem eager to discuss and debate extensively.

To tackle the challenges coming from both data quantity and data complexity, social network analysts must adopt and adapt AI-powered tools to increase the speed and scale of their work, and to reduce the natural noise among the collected data. We presented a method for exploring a social media platform through its communities, to qualify their social and topical cohesion, and to illustrate their dispersion around their main topics of discussion.

Glossary

This is a brief glossary of some of the vocabulary employed in graph theory for readers who may not be familiar with the discipline.

Node **(or vertex)**	Basic, fundamental unit. A node usually represents an entity (e.g. a person, company, place, ship, bank account…)	
Edge **(or link)**	The connection between two nodes. It can be directed (represented by an arrow), or undirected. It often carries a weight (the strength of the relation) or more information.	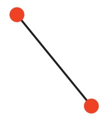
Graph **(or network)**	The object formed by two sets: *V* a set of vertices, and *E* a set of edges between (some of) these vertices.	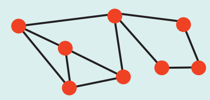
Community	A set of nodes that are more closely linked together, with relatively fewer links towards the exterior of the community.	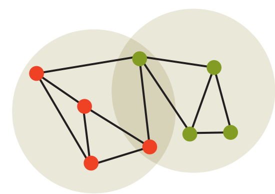
TPR	The proportion of nodes that are linked to two other nodes and linked together themselves. On the example, the red group has a TPR of 1 as each node is part of a group of 'friends'; the green group has a TPR of 0.	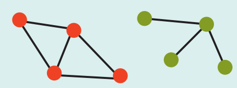
Modularity	A quality measure of communities, based on the number of external (outgoing) links, in dark, as opposed to internal edges (dashed)	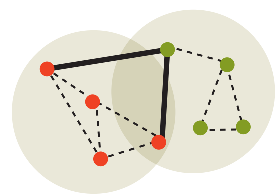
Degree	The degree of a node is the number of edges coming or going from it. When edges are directed, a distinction is made between in-degree (incoming edges) and out-degree.	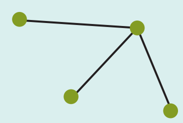

Endnotes

1 J. Shaheen, Network of Terror: How Daesh Uses Adaptive Social Networks to Spread its Message (Riga, Latvia: NATO Strategic Communications Centre of Excellence, 2015).

2 Rita Katz, 'A Growing Frontier for Terrorist Groups: Unsuspecting Chat Apps', Wired, 1 September 2019.

3 Majid Alfifi, Parisa Kaghazgaran, James Caverlee, and Fred Morstatter, 'Measuring the Impact of ISIS Social Media Strategy', 2018.

4 Mirom Lakomy, 'Cracks in the Online "Caliphate": How the Islamic State is Losing Ground in the Battle for Cyberspace', Perspectives on Terrorism Vol 11 No 3 (2017): 2–15.

5 A. T. Chatfield, C. G. Reddick, and U. Brajawidagda, 'Tweeting Propaganda, Radicalization and Recruitment: Islamic State Supporters Multi-sided Twitter Networks' in Proceedings of the 16th Annual International Conference on Digital Government Research, pp. 239–49. ACM, 2015.

6 Val Powell, 'Anonymous Upstaged By 'Ghost Security Group' In Cyber War Against ISIS', Inquisitr, 19 November 2015.

7 Humanshuddle, 'ISIS: Sunset on the "decline narrative"', Online Jihad, 1 June 2018.

8 Elian Peltier, 'Fabien Clain, Prominent French Voice of ISIS, Is Reported Killed in Syria', The New York Times, 28 February 2019.

9 Vincent D. Blondel, Jean-Loupe Guillaume, Renaud Lambiotte, and Etienne Lefebvre, 'Fast Unfolding of Communities in Large Networks', Journal of Statistical Mechanics: Theory and Experiment 2008 (10): P10008.

10 K. H. Lim and A. Datta, 'Finding Twitter Communities with Common Interests Using Following Links of Celebrities' in Proceedings of the 3rd International Workshop on Modelling Social Media, pp. 25–32. ACM, 2012.

11 G. Gadek, A. Pauchet, N. Malandain, K. Khelif, L. Vercouter, and S. Brunessaux, 'Measures for Topical Cohesion of User Communities on Twitter' in Proceedings of the International Conference on Web Intelligence, pp. 211–18, ACM, 2017.

12 Discord app website. Discord's servers allow users to communicate using the voice tool or chat in a channel (discussion room). Discord also provides a direct messaging option, but these communications cannot be seen by anyone else. Anyone can read a channel that is set to 'public'; most Discord servers are set to 'private'.

13 Reddit homepage.

14 Twitter homepage. The search feature on Twitter is facilitated by the use of hashtags: author-signalled keywords identified by the '#' symbol. Twitter data collection can be performed through either the Search API, requesting several hundred tweets a t a time; or through the Stream API, posting a query to receive future tweets under a keyword presence condition.

15 John L. Esposito (ed.), entry 'Fard al-Kifayah' in The Oxford Dictionary of Islam (Oxford University Press, 2003), also in Oxford Islamic Studies Online.

16 R. Lara-Cabrera, A. G. Pardo, K. Benouaret, N. Faci, D. Benslimane, and D. Camacho, 'Measuring the Radicalisation Risk in Social Networks', IEEE Access Vol 5 (2017) : 10892–10900.

17 Fifth Tribe, 'How ISIS Uses Twitter', kaggle.com, no date.

18 A. Bermingham, M. Conway, L. McInerney, N. O'Hare, and A. F. Smeaton, 'Combining Social Network Analysis and Sentiment Analysis to Explore the Potential for Online Radicalisation' in ASONAM 2009—Advances in Social Networks Analysis and Mining, 20–22 July 2009, Athens, Greece, pp. 231–36.

19 Pennebaker, J. W., Chung, C. K., et al., 'Computerized Text Analysis of Al-Qaeda Transcripts' in K. Krippendorff & M. Bock (eds) A Content Analysis Reader (Thousand Oaks, CA: Sage, 2008).

20 B. McFarlane, 'Online Violent Radicalisation (OVeR): Challenges Facing Law Enforcement Agencies and Policy Stakeholders' in ARC Linkage Project on Radicalisation—Conference, 2010.

21 M. Ahmed and F. L. George, 'A War of Keywords: How Extremists are Exploiting the Internet and What To Do About It', (Center on Religion and Geopolitics, 2016).

22 Blondel et al, 'Fast Unfolding of Communities'.

23 A. Lancichinetti, F. Radicchi, J. J. Ramasco, and S. Fortunato, 'Finding Statistically Significant Communities in Networks', *PLOSONE* 6(4): e18961, 2011.

24 Yong-Yeol Ahn, James P Bagrow, and Sune Lehmann, Letter 'Link Communities Reveal Multiscale Complexity in Networks', Nature Vol. 466 No. 7307 (2010): 761–64.

25 Gennaro Cordasco and Luisa Gargano, 'Community Detection via Semi-synchronous Label Propagation Algorithms' in 2010 IEEE International Workshop on Business Applications of Social Network Analysis (Bangalore, India: IEEE, 2010) pp. 35–42.

26 Steve Gregory, 'Finding Overlapping Communities in Networks by Label Propagation', New Journal of Physics Vol 12, October 2010: 103018.

27 Quoc Le and Tomas Mikolov, 'Distributed Representations of Sentences and Documents' in Proceedings of the 31st International Conference on Machine Learning, PMLR 32 (2): 1188–96, 2014.

28 Thomas K. Landauer, Peter W. Foltz, Darrell Laham, 'An introduction to latent semantic analysis', Discourse Processes Vol 25 No 2–3 (1998): 259–84.

29 Zee Media Bureau, 'Terrorist Maulana Masood Azhar, Head of Pakistan-backed Jaish-e-Mohammad, Dead: Sources', Zee News, 3 March 2109.

30 Rukmini Callimachi, 'ISIS Caliphate Crumbles as Last Village in Syria Falls', The New York Times, 23 March 2109.

31 'ISIS Bride Vows Son "Will Grow Up to Be a Jihadist" as Crumbling Caliphate Evacuated', Breitbart London, 23 February 2019.

32 Joseph Shasheen, Detrimental Use of Social Media and the Case of Daesh: An Information Warfare Perspective, Riga Strategic Communication Dialogue: Perception Matters, August 2015.

33 Geoffrey L. Cohen, 'Party over policy: The dominating impact of group influence on political beliefs', Journal of Personality and Social Psychology Vol. 85, Issue 5 (2003): 808.

34 L. J. P. van der Maaten and G.E. Hinton, 'Visualizing Data Using t-SNE', Journal of Machine Learning Research 9 (November 2008): 2579–605.

Prepared and published by the

NATO STRATEGIC COMMUNICATIONS CENTRE OF EXCELLENCE

The NATO Strategic Communications Centre of Excellence (NATO StratCom COE) is a NATO accredited multi-national organisation that conducts research, publishes studies, and provides strategic communications training for government and military personnel. Our mission is to make a positive contribution to Alliance's understanding of strategic communications and to facilitate accurate, appropriate, and timely communication among its members as objectives and roles emerge and evolve in the rapidly changing information environment.

Operating since 2014, we have carried out significant research enhancing NATO nations' situational awareness of the information environment and have contributed to exercises and trainings with subject matter expertise.

www.stratcomcoe.org | @stratcomcoe | info@stratcomcoe.org