

# Combining Social Network Analysis and Sentiment Analysis to Explore the Potential for Online Radicalisation

Adam Bermingham<sup>1</sup>, Maura Conway<sup>2</sup>, Lisa McInerney<sup>2</sup>,  
Neil O’Hare<sup>1</sup>, Alan F. Smeaton<sup>1</sup>

<sup>1</sup>CLARITY: Centre for Sensor Web Technologies  
and <sup>2</sup>School of Law and Government,  
Dublin City University, Glasnevin, Dublin 9, Ireland.

## Abstract

*The increased online presence of jihadists has raised the possibility of individuals being radicalised via the Internet. To date, the study of violent radicalisation has focused on dedicated jihadist websites and forums. This may not be the ideal starting point for such research, as participants in these venues may be described as “already made-up minds”. Crawling a global social networking platform, such as YouTube, on the other hand, has the potential to unearth content and interaction aimed at radicalisation of those with little or no apparent prior interest in violent jihadism. This research explores whether such an approach is indeed fruitful. We collected a large dataset from a group within YouTube that we identified as potentially having a radicalising agenda. We analysed this data using social network analysis and sentiment analysis tools, examining the topics discussed and what the sentiment polarity (positive or negative) is towards these topics. In particular, we focus on gender differences in this group of users, suggesting most extreme and less tolerant views among female users.*

## 1. Introduction

Jihadists have significantly increased their online presence since 9/11 [16] raising the possibility of exposure to online violent jihadist content resulting in individual users being radicalised via the Internet. Online radicalisation is here conceived as a process whereby individuals, through their online interactions and exposure to various types of Internet content, come to view violence as a legitimate method of solving social and political conflicts. Recent approaches to the study of violent radicalisation have focused on dedicated jihadist websites and forums. We suggest that this may not be the ideal starting point for such research, however, as participants in dedicated online jihadist venues may

be described as “already made-up minds”. In crawling YouTube<sup>1</sup> we can look at interaction aimed at those with little or no prior interest in violent jihadism. The present paper builds on previous research on the links between jihadi video and online radicalisation, and the contribution is a detailed analysis of a real YouTube dataset. This analysis uses an application of sentiment-, lexical- and social network analysis, which allows us to examine and characterise the users of radicalised forums, with particular emphasis on gender-based differences between users.

## 2. Related Work

In the sub-sections below we introduce related research into online radicalisation, followed by a summary of a previous case study, in which we identified the YouTube group that we focus on in this paper. We also describe related work in the field of sentiment analysis.

### 2.1. Online Radicalisation

A number of individuals and groups are reported to have used the Internet extensively, both in the process of radicalisation and, in some instances, attack planning and preparation. Hussain Osman, one of the London bombers, claimed to have been influenced by watching Internet video footage of the Iraq conflict and reading about jihad online [2], and the perpetrators of the 2005 Khan al-Khalili bombing in Cairo downloaded bomb-making instructions from a jihadist website [16]. There is broad agreement amongst both researchers and policymakers that the Internet is a venue for jihadi radicalisation although there is a dearth of empirical academic research addressing this issue [15],[5].

Marc Sageman claims to redress this in his *Leaderless Jihad* [16], probably the most well known text to explicitly address online radicalisation. He has been criticised by

<sup>1</sup><http://www.youtube.com>

Bruce Hoffman [8], however, precisely for failing to employ social network theory and associated methods to evidence his argument as to al-Qaeda's present structure and operation, including its Internet component.

In previous research we explored remedying some of the deficiencies in the online radicalisation literature identified by Hoffman, in particular running social network analysis on our data. Where we disagree with Sageman, and others, however is in our belief that the Internet is not neatly divided in two with information provision (i.e. passive websites) on the one side and interactivity (i.e. forums, chat rooms, etc.) on the other. YouTube, for example, seamlessly blends these two components and much more through video- and text-based information provision and interactivity through users commenting on videos, writing on each others' profiles, selecting friends, joining groups, etc. While traditional forums and chat rooms are interactive, the type of interactivity facilitated falls short of the type of interactivity possible on social networking sites and, in particular, YouTube with its multi-modal nature.

So, while Sageman and others focus on jihadist forums as the likeliest venues for radicalisation, the present research focuses on Internet users with no apparent prior links to jihadism who may have the potential to be drawn into radical circles via their interactions on a popular global site.

## 2.2. Our Previous Online Radicalisation Analysis

This paper builds on previous research into the online supporters of jihad-promoting video content on YouTube [6]. In the study, videos relevant to martyrdom were identified and discourse and content analysis techniques were employed on the associated comments and user profiles to classify individuals as 'supporter', 'critic', 'neither' or 'impossible to determine'. Demographic details were manually extracted from all user profiles, including age and location. Additional details were obtained manually from text contained within user profiles (e.g. national or ethnic markers).

A number of users potentially involved in radicalisation processes were identified and their interactions were detailed in a series of mini-case studies. Two users in particular that we identified were both members of the same 600+ (now 700+) member YouTube group the stated purpose of which was the "conversion of infidels". For this paper we have used this group as the basis for our analysis; we have not provided exact figures for group membership in an effort to prevent easy de-anonimisation of our data [11]. The case studies suggested that examining the similarities and differences between the male and female group members would also prove a valuable exercise. An increased role for women in the violent jihadist movement arising out of their ability to involve themselves online has been argued by Sageman and others, and in this paper we profile male

and female users using quantitative methods.

## 2.3. Sentiment Analysis

Sentiment analysis is the task of identifying positive and negative opinions, emotions, and evaluations [18]. Sentiment analysis has been used with much success in fields where users have an obvious subjective agenda such as movie reviews [14] or blogs [13]. It is much less clear how sentiment analysis techniques can be employed in the context of social network analysis where the language tends to be more freeform and informal.

We propose a novel application of a state-of-the-art sentiment analysis technique to examine social relationships and networks in the context of radicalization. In this paper we use a dictionary-based polarity scoring method to assign positivity and negativity scores to YouTube profiles and comments. This enables us to characterise users and groups of users by their sentiment towards a set of concepts found in our test corpus that are deemed to be of particular interest to jihadists.

## 3. Methodology

By analysing the YouTube user profile information and discussion initiated by members of the group identified, we sought to answer the following two research questions:

- Was this group populated by radicals who were in a position to draw others into their sphere of influence ?
- What were the differences, if any, in terms of radicalness between the content posted, and the interactions engaged in, by the male and female group members ?

Data collection and analysis consisted of a number of steps: a YouTube crawl to gather relevant data, a network analysis of this data, and lexical analysis of the corpus to inform the sentiment analysis of the documents gathered.

### 3.1. YouTube Crawl

We crawled all comments and user profiles on the YouTube group discussed in Section 2.2 above, in order to gather real world data from a forum that we believed might have a radicalising agenda. The group, the name and other details of which we are omitting for privacy reasons, today has 700+ members. An automated crawl of the group's YouTube data was first conducted on 8 January, 2009 and again on 16 January, 2009.

The initial crawl gathered user comments and the latter gathered user profile information. Comments were read from the "Channel Comments" portion of group members'

	Users	Comments	Total
Total	13,700+	122,016	135,000+
Group Members	700+	21,587	22,000+
Non Members	13,000+	100,416	113,000+
Males	5,967	40,798	10,765
Female	2,715	21,120	23,835
Unknown Gender	5,081	60,098	65,179

**Table 1. User profiles and comments in our corpus. Exact group members omitted for privacy reasons.**

profiles. User profile information was read for all group members and for any YouTube user who had posted on a group member’s channel. The corpus used in this paper is thus comprised of both user comments and profile text. In total there are over 13,000 profiles and 122,011 comments, giving 135,000+ unique ‘documents’, each either a user profile or a user comment. More details of the corpus are given in Table 1. The 700+ group members were responsible for over 22,000 documents while 13,000+ other users were responsible for over 113,000 documents. Of the 8,682 users who declared their gender, 2,715 (31%) are female. We do not have gender information available for over 5,000 users. The proclaimed age of group members ranged from 14 to 107 years. We removed one outlier of 107 years in advance of reporting age data statistics in Table 2, where we see that females are on average younger than males in this corpus.

	Min (years)	Max (years)	Mean (years)
All	14	56	24.5
Male	14	56	25.2
Female	15	41	22.8

**Table 2. Age profile of the users in our corpus**

### 3.2. Sentiment Analysis

Sentiment Analysis was performed on all of the documents gathered during the crawl. Initially, we ran a system we developed for our participation in TREC 2008, which focused on opinion detection and opinion polarity detection in a large scale blog corpus [3]. Differences in the nature of the data, however, made the results unreliable. In YouTube, most comments are considerably less than 50 words in length, much shorter than the typical blog length in the TREC corpus of hundreds of sentences. Another major difference was that there was little evidence of subjectivity

in the YouTube text. Often, when a YouTube user expresses an opinion they simply state it rather than qualifying it with “I think...” or “I feel...”. This behaviour is not seen in the blog corpus where authors are keen to distinguish opinion from fact in their posts. All of this meant that the components in the TREC system designed to detect subjectivity were not useful in the context of the YouTube data, and so hindered performance.

For these reasons, we decided to use only the polarity module (“Lexicon Module”) from our system to score documents for positivity and negativity. This module uses a lexicon, SentiWordNet, which assigns positivity and negativity scores to synset entries in WordNet [7]. Two improvements were made to the TREC algorithm:

- Part of speech was used to help disambiguate between SentiWordnet entries. We used the Stanford Maximum Entropy Part of Speech Tagger and converted the Penn Tree Bank parts of speech to WordNet parts of speech [17].
- RitaWN was used to employ WordNet stemming before looking up terms in SentiWordNet [1].

As the polarity orientations of the terms do not necessarily correspond to subjectivity, a document may simply discuss an aspect of a concept using polarised terms, but without expressing opinion. For example, the sentence “He is sick and tired today”, has two negatively oriented terms, “sick” and “tired”, yet it remains a negative statement of fact rather than a negative opinion.

Running sentiment analysis on a document generates a positive score and a negative score. This corresponds to the mean term orientation in the document. For profiling genders and concepts we filtered documents by author gender and concept terms respectively. For a set of documents, the positivity sentiment score is defined as the mean positivity score for the documents in that set; negativity scores are calculated similarly.

## 4. Findings and Discussion

In the subsections below we present a lexical analysis of the crawled corpus, followed by sentiment analysis results and social network analysis results.

### 4.1. Lexical Analysis

In order to extract information about the lexical features of the corpus, all comments and user profiles were indexed by the Terrier search engine [12]. Stopwords were removed and Porter’s stemming algorithm was used to remove morphological variants of words. The search engine index was used to extract various statistics about the corpus lexicon, in particular information about frequencies of word usage

over various sub-sets of the corpus. Tables 3, 4 and 5 outline the most frequently used terms for all users, males and females respectively. In each table, the following metrics are reported:

- Term frequency (TF), the number of times the term occurs in the entire collection.
- Document frequency (DF), the number of unique comments or user profiles that the term occurs in.
- User frequency (UF), the number of unique users who use this term.

Term	TF	DF	UF
allah	32,777	19,376	3,991
islam	23,003	14,596	3,174
jesus	22,266	6,045	687
like	21,977	13,427	3,441
muslim	16,431	12,355	3,231
who	16,209	9,039	2,602
god	16,125	9,981	1,851
thank	15,507	13,760	3,603
will	14,552	10,432	3,017
salam	13,239	11,558	2,215

**Table 3. Top 10 most frequently occurring terms for all users**

As one would expect in such a forum, the documents in this collection are dominated by terms relating to religion. Table 1 shows that the top 10 most used terms include 5 that exclusively relate to religion with the two most frequently used terms being *Allah* and *Islam*. Interestingly, discussion is not restricted to Islam, and *Jesus* is the third most frequently used term, although it appears in relatively few unique documents from relatively few unique users.

Term	TF	DF	UF
jesus	17,716	3,659	294
like	11,239	4,717	1,294
who	10,130	4,563	977
allah	8,840	5,867	1,549
islam	8,002	5,476	1,297
muslim	7,237	5,431	1,341
follow	6,773	2,943	295
god	5,613	3,525	751
will	5,512	3,991	1,112
peace	5,254	3,833	676

**Table 4. Top 10 most frequently occurring terms for male users**

Comparing male and female users, it is striking that *Jesus* is the most frequently used term by males, although it does not register at all in the top 10 for females. The low user frequency of *Jesus* for males suggests, however, that this comes from a relatively small number of users who use this term very frequently. It is also noteworthy that the only exclusively religious term in the top 10 for females is *Allah*.


Term	TF	DF	UF
allah	4,726	3,192	846
salam	3,842	2,732	560
sis	3,166	2,923	382
ur	2,976	2,199	519
like	2,712	1,664	724
well	2,563	2,451	427
am	2,514	1,995	575
lol	2,492	2,015	359
thank	2,345	2,244	828
love	2,284	1,516	791

**Table 5. Top 10 most frequently occurring terms for female users**


## 4.2 Sentiment Analysis Results

Lexical analysis of the corpus allowed us to determine what the users in the group were talking about, while sentiment analysis techniques were able to inform us about the opinions or attitudes of the users towards these topics. To determine target topics for which we could extract the user sentiment, we constructed ‘concepts’ of potential interest to jihadists from the top 50 most frequently used terms. These 8 concepts were *America, Christianity, Islam, Israel, Judaism, Mubarak, Palestine, al-Qaeda*. To find all documents (ie. comments or user profiles) relevant to these concepts, each concept was expanded with spelling variants, synonyms etc. For example, for *Christianity* all documents containing any of the following terms (among others) were considered relevant: *Jesus, cross, Christian, Bible*.

The results of the sentiment analysis for each of these concepts are presented in Figures 1 and 2. Figure 1 shows the sentiment analysis results for males. We can see that for certain topics, such as *America, Christianity* and *Palestine*, the level of positivity and negativity are broadly similar. The most overwhelmingly positive topic for males is *Mubarak*, followed by *Islam*. Interestingly, *Judaism* has higher positivity than negativity, while the opposite can be said for the topic *Israel*, suggesting that many of these males are tolerant of the Jewish religion while opposing the Israeli state. The very high positive sentiment for Egyptian president Hosni Mubarak is noteworthy, although we cannot explain this without further exploration and analysis.



**Figure 1. Sentiment Analysis Results: Males**



**Figure 2. Sentiment Analysis Results: Females**

Figure 2 shows the sentiment analysis results for females. Again, the most positive topics are *Mubarak* and *Islam*. What is most striking, in terms of the difference between genders, is greater female positivity towards the topic *al-Qaeda* and the greater female negativity towards the topic *Judaism*. This suggests a more extreme viewpoint among the female users in the corpus: unlike the men, they seem unwilling to distinguish between the Israeli state and the Jewish religion. They are also much more positive towards al-Qaeda. Negative female sentiment towards Christianity is strong too, again suggesting a greater lack of tolerance of other religions. Indeed, if the difference between the positivity and negativity score for each concept is considered for each gender, the only reversal of polarity is seen in the non-Islam religious concepts, Christianity and Judaism. For these concepts, males are more positive than they are negative and females are markedly more negative than they are positive.

### 4.3 Social Network Analysis

Centrality relates to the extent that an actor is central to a network. For example, an actor on the fringe of a network with only one connection would score low on centrality measures as compared to an actor that connects with

all actors within the network. There are various centrality measures used in network analysis, with *betweenness* being a typical measure used to identify actors who are potential influencers. Betweenness is defined as the “extent to which an actor mediates, or falls between any two actors on the shortest path between those actors”. This is usually averaged across all possible pairs in the network [9]. Carley *et al.* [4] went on to develop a cognitive demand measure, the function of which is to find other key emergent actors that may not feature in the top ranking based on betweenness centrality measures. We decided to use both of these measures to identify key actors.

The top ten ranking group members in betweenness centrality scores featured two females (ranked 2nd and 8th), while the remaining eight members had not declared their gender on YouTube. The top ten ranking group members in cognitive demand scores revealed a largely similar set of actors, the gender of eight of whom is unknown: two females ranking highly (1st and 8th). This places a female in the position of overall emergent leader. What is also interesting, is that the individual ranked overall emergent leader is also ranked second in betweenness centrality scores. She is therefore also strongly positioned as a potential influencer. The two females who rank eighth in both measures are different users. Of the top ten users who frequently rank in the top three of various social network measures, five are of unknown gender and five identify themselves as female. It is quite possible that at least some of the five users of unknown gender are in fact female because of a masculinist bias apparent in the Muslim world that might cause women to cloak their female identity. For example, jihadist forums are often segregated on the basis of gender but it is thought that many females are populating such forums as it is difficult to validate a user’s gender online [16].

With regard to network characteristics, both network density and average communication speed results were obtained. Network density is the ratio of the number of actual links to the number of possible links. Average communication speed is the average speed with which any two actors can interact, calculated as the inverse of the shortest path lengths between actor pairs. As expected, and in line with previous work reporting on network characteristic differences amongst males and females [10], females scored higher in terms of both network density (0.0092 for females, 0.0022 for males) and average communication speed (0.3584 for females, 0.2282 for males). While such results may be expected, it is important not to disregard this finding, in conjunction with those above, as it indicates a potentially increased leadership role for women online than they would generally be held to have within jihadi circles.

## 5. Conclusions and Future Work

In this paper we presented an analysis of online social network content in the context of jihadi radicalisation. We have profiled users and interactions in a targeted YouTube group using an application of sentiment-, lexical- and social network analysis techniques. In terms of our research questions, this group has no core radicalising function and is mostly devoted to religious discussion, particularly Islam.

With regard to our second question however, we have been able to make a number of interesting observations about the differences in the nature of the discussion and interactions as between the male and female members of the group. Of particular note is the tendency of the women to be more sympathetic to political violence actors and also not to distinguish between the Jewish religion and the state of Israel, a division that appears to be clearly recognised by male group members. Another finding, is that the higher one's status within the group, on several different measures, the less likely one is to provide information about one's gender. Does this mean, for example, that there are more women active at the highest levels within the group than revealed here?

The data gathered for this work is very focused on the targeted YouTube group. We would like to expand the breadth of the corpus allowing us to build up a more complete social network. We also plan to integrate temporal information to allow us to better examine the changing character of the network and to analyse reaction to events e.g. the posting of a new video, real-world news events. With this extended corpus we can more fully explore the radicalising processes. In terms of sentiment analysis and lexical analysis, a larger corpus would also be desirable. For concepts outside the 50 most frequently used terms, and certainly below 100, there are few relevant documents. The concept development phase was therefore restricted to common terms, whereas looking at rarer but more intense concepts could prove more valuable. The lexicon used for sentiment is not domain specific to our purposes and thus is subject to the problems of polysemy and synonymy. In future work we will mitigate against this by using supervised learning or other approaches to better capture the unique language of this domain. There are many examples of non-English text in the corpus also, particularly Arabic, which we feel contains valuable sentiment on behalf of the users. In future work we will integrate the non-English text into our sentiment analysis algorithm.

Finally, while research on online radicalisation has to date focused upon violent jihadis and their activities, extending such research to neo-Nazis or even animal rights groups with a history of violence might also prove useful, while also initiating the possibility for comparative cross-ideological research in the area.

## Acknowledgments

This work is supported by Science Foundation Ireland under grant 07/CE/I1147, and by Enterprise Ireland under grant IP/2008/0549.

## References

- [1] <http://www.rednoise.org/rita/wordnet>.
- [2] A. N. Awan. Radicalization on the Internet? The virtual propagation of jihadist media and its effects. *Journal of the Royal United Services Institute*, 152(3), June 2007.
- [3] A. Bermingham, A. Smeaton, J. Foster, and D. Hogan. DCU at the TREC 2008 Blog Track. In *The Seventeenth Text REtrieval Conference (TREC 2008) Proceedings*. NIST, 2008.
- [4] K. M. Carley, M. Dombroski, M. Tsvetovat, J. Reminga, and N. Kamneva. Destabilizing dynamic covert networks. In *8th International Command and Control Research and Technology Symposium*, pages 726–735, 2003.
- [5] F. Cilluffo, G. Saathoff, S. Cardash, J. Lane, and A. Whitehead. Networked radicalization. Technical report, George Washington University and University of Virginia, 2007.
- [6] M. Conway and L. McInerney. Jihadi video and auto-radicalisation: Evidence from an exploratory YouTube study. In D. Ortiz-Arroyo, H. Larsen, D. Zeng, D. Hicks, and G. Wagner, editors, *Intelligence and Security Informatics*, pages 108–118, 2008.
- [7] A. Esuli and F. Sebastiani. Sentiwordnet: A publicly available lexical resource for opinion mining. pages 417–422, 2006.
- [8] B. Hoffman. The myth of grass-roots terrorism. *Foreign Affairs*, Vol 87, 2008.
- [9] P. R. Monge and N. S. Contractor. *Theories of Communication Networks*. Oxford University Press, US, 2003.
- [10] G. Moore. Structural determinants of men's and women's personal networks. *American Sociological Review*, pages 726–735, 1990.
- [11] A. Narayanan and V. Shmatikov. De-anonymizing social networks. (arXiv:0903.3276), Mar 2009.
- [12] I. Ounis, C. Lioma, C. Macdonald, and V. Plachouras. Research directions in terrier: a search engine for advanced retrieval on the web. *Novatica/UPGRADE Special Issue on Web Information Access, Ricardo Baeza-Yates et al. (Eds)*, 2007.
- [13] I. Ounis, C. MacDonald, and I. Soboroff. Overview of the TREC-2008 Blog Track. In *The Seventeenth Text REtrieval Conference (TREC 2008) Proceedings*. NIST, 2008.
- [14] B. Pang, L. Lee, and S. Vaithyanathan. Thumbs up? Sentiment classification using machine learning techniques. *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 79–86, 2002.
- [15] J. Ryan. *Countering Militant Islamist Radicalisation on the Internet*. Institute of European Affairs, Dublin, 2007.
- [16] M. Sageman. *Leaderless Jihad*. University of Pennsylvania Press, Philadelphia, 2008.

- [17] K. Toutanova and C. Manning. Enriching the knowledge sources used in a maximum entropy part-of-speech tagger. *Proceedings of the 2000 Joint SIGDAT Conference EMNLP/NLC, 63-71, 2000*, pages 63–70, October 2000.
- [18] T. Wilson, J. Wiebe, and P. Hoffmann. Recognizing contextual polarity in phrase-level sentiment analysis. *Proceedings of the 2005 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 347–354, 2005.