# Mining YouTube to Discover Extremist Videos, Users and Hidden Communities

Ashish Sureka, Ponnurangam Kumaraguru, Atul Goyal, and Sidharth Chhabra*

Indraprastha Institute of Information Technology, New Delhi - 110078
{ashish,pk,atul08015}@iiitd.ac.in, sidharth.chhabra2011@coe.dce.edu
http://www.iiitd.edu.in/~pk/precog/

**Abstract.** We describe a semi-automated system to assist law enforcement and intelligence agencies dealing with cyber-crime related to promotion of hate and radicalization on the Internet. The focus of this work is on mining YouTube to discover hate videos, users and virtual hidden communities. Finding precise information on YouTube is a challenging task because of the huge size of the YouTube repository and a large subscriber base. We present a solution based on data mining and social network analysis (using a variety of relationships such as friends, subscriptions, favorites and related videos) to aid an analyst in discovering insightful and actionable information. Furthermore, we performed a systematic study of the features and properties of the data and hidden social networks which has implications in understanding extremism on Internet. We take a case study based approach and perform empirical validation of the proposed hypothesis. Our approach succeeded in finding hate videos which were validated manually.

**Keywords:** Information Retrieval, Hate and Extremism Detection, Security Informatics, YouTube Content Analysis, Social Network Analysis.

## 1 Introduction

Video-sharing websites such as YouTube have become a channel for spreading extremism and being used as a Internet based distribution platform for like-minded people to interact, publicize and share their ideologies [5]. Due to low publication barrier (self-publishing model) and anonymity, websites such as YouTube contains a large database of user generated content (UGC) in the form of videos and textual comments which are malicious and racist (this is despite several efforts by YouTube administrators to remove offensive content based on users complaints) [1], [6]. Online extremism and hate content can have a negative impact to the society and the prevalence of such easily accessible content (as

* The author is a student at DTU (Delhi Technological University). However, this work was performed while the author was doing his internship at IIIT-D (Indraprastha Institute of Information Technology, Delhi).

P.-J. Cheng et al. (Eds.): AIRS 2010, LNCS 6458, pp. 13–24, 2010.

anyone can watch online videos and does not even need to create an account) is thus a major concern to the people, government and law enforcement agencies. Solutions to counter cyber-crime related to promotion of hate and radicalization on the Internet is an area which has recently attracted a lot of research attention. The information need of a law enforcement agent or a security analyst detecting cyber-hate on YouTube (the focus of this work) is the following[1]:

1. *Videos* promoting hate and extremism
2. Influential *users* and *leaders* playing a central role in spreading such sentiments
3. Virtual *communities* and hidden social networks of people with the shared agenda and interest

Furthermore, a study of the *properties* of hate related YouTube videos, users and communities can lead to a better understanding of the problem and have implications in designing solutions to address such a problem. Finding extremist videos, users and virtual communities on YouTube is a technically challenging task due to the vastness of YouTube repository in terms of the number of videos, users and the different types of relationships between them. Keeping in mind the need to devise solutions for countering and studying cyber-hate, the *research aim* of this work is:

1. To investigate solutions to support a security analyst to extract actionable information from YouTube with respect to cyber-hate and extremism
2. To investigate the properties and features of the extremist content, users and hidden communities on YouTube

We now compare and contrast our work from closely related previous research and in context to related work list the unique contributions of this paper.

**Analysis of Online Hate Videos.** Reid et al. studied extremist and terrorist groups' videos and perform a content analysis of 60 jihadi videos. They analyze attributes like video types (documentary, suicide attack, propaganda, instruction), production features (special effects, subtitles) and communication approaches (audience segmentation) [11]. Adam et al. study online radicalization by analyzing dataset from a group within YouTube. They studied user-profile information, perform sentiment and lexical analysis of forum comments, apply social network analysis and derive insights on gender differences in views around jihad-promoting content on YouTube [1]. Conway et al. perform an analysis of jihadi video content on YouTube with a focus on martyr-promoting material from Iraq [6]. They studied a sample of 50 videos uploaded by 30 individual users and analyzed user profiles (categorizing users as supporters or critic), comments (total of 1443 comments by 940 separate users), demographic details (age and current location), popularity metrics (such as number of views, comments and ratings).

[1] Based on inputs from senior officers from law enforcement and intelligence agencies.

**Analysis of Online Hate Blogs.** Chau et al. present a semi-automated approach to analyze virtual hate communities in blogosphere. They analyze anti-Blacks hate groups and bloggers on Xanga which is a popular blog hosting website [4]. The similarities between the work by Chau et al. and this paper are the application of network analysis and text analytics to analyze subscription linkages and textual comments respectively. While the motivation between the two works is the same, Chau et al. analyze anti-Blacks hate groups in Blogosphere whereas we study anti-India hate groups in YouTube.

**Analysis of YouTube Social Network.** There are several papers on the study of YouTube video sharing community and social network analysis. Due to the limited space in the paper, we discuss a few recent and closely related work. Biel et al. study the properties and structure of YouTube social networks with a focus on analyzing the network of subscriptions (large-scale static and dynamic analysis) [2]. Santos et al. collect a representative sample of YouTube using a crawler and analyze the structural properties and social relation-ships among users, among videos, and between users and videos [12]. Mislove et al. examine data gathered from YouTube Flickr, YouTube, LiveJournal, and Orkut and presents a large-scale measurement study and analysis of the structure of multiple online social networks [9].

**Research Contributions.** This study is an attempt to advance the state-of-the-art in the area of *cyber-hate analysis and detection*. The study focuses on YouTube online video sharing and social networking website. In context to the related work, the specific novel contributions of this paper are:

1. A *general framework* to facilitate security analysts and intelligence agencies to identify hate and extremist content, users and hidden communities on *YouTube*. To the best of our knowledge, this paper is the first study to perform an *integrated analysis of a wide variety of user and video attributes and relationship in the context of cyber-hate in YouTube*. The study investigates popularity metrics, user and video features, network relationships such as friends, favorites/playlists and subscriptions, related video relationship and performs linguistic analysis of user comments.
2. A method to discover *hate content, users and communities* on YouTube by leveraging a variety of user-user (e.g. friends and subscriptions) and user-video (e.g. uploader and favorites/playlists) relationship using social network analysis tools and techniques.
3. We believe that it is important to study cyber-hate having a focus on different nations, religions and communities so that comparisons and references can be made to better understand the problem from different perspectives. To the best of our knowledge, this is the first *India-centric* academic research on analyzing cyber hate on a video sharing and social networking website.

The remainder of the paper is organized as follows: In Section 2, we discuss the framework that we developed to analyze the data set from YouTube, the methodology by which we collected the data set from YouTube, and the results from the data analysis that we performed on the collected data. Finally, in Section 3, we conclude the paper with usefulness of this research work.

# 2 Empirical Analysis

## 2.1 Proposed Framework

Figure 1 presents the proposed framework to analyze YouTube repository for extracting hate-promoting users, videos and communities. As shown in Figure 1, the starting point is a seed list of videos (generated manually) which is then used as a seed reference to extract more users and videos (through a process called as *bootstrapping* or *snow ball sampling* [10]). Figure 1 shows various network analysis and linguistic analysis modules to analyze several types of user and video relationships present in YouTube. Each of the component (user comment analysis using natural language processing techniques, socio-centric and ego-centric graph analysis for community discovery, network analysis based on friends and subscription relationship for uncovering additional like-minded users) illustrated in the Figure 1 is described in the following sections of the paper. As shown in Figure 1, the ultimate goal (output produced by the system) is to assist a security analyst in retrieving and visualizing relevant useful and actionable information.

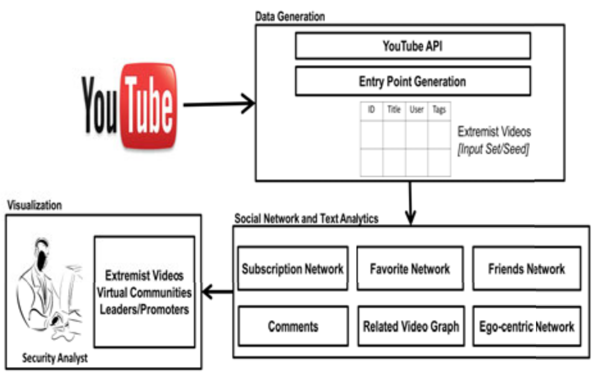


**Fig. 1.** The framework that we used in the analysis. We used the YouTube API to download the details of the videos, and users. We stored it in a database and ran our analytics tool on it to produce statistics and visualization.

## 2.2 Seed Dataset (Entry-Point)

The approach presented in this paper assumes that an entry point (seed list) is provided to the system as an input which is used as a base to uncover additional videos, users and communities. The authors of this paper with the support of six undergraduate students identified 75 YouTube videos based on a manual search. We created a list of Video IDs and the values of all other fields are

**Table 1.** Illustrative list of videos and selected popularity indicators in the input set (UD: Upload Date (all videos were uploaded in 2009), DS: Duration in Seconds, AR: Average Rating, NR: Number of Raters, VC: View Counts, NF: Number of Favorites, NC: Number of Comments). We have anonymized the titles of the videos, user IDs, names, and video IDs in the paper.

| **Title** | **Category** | **UD** | **DS** | **AR** | **NR** | **VC** | **NF** | **NC** |
|---|---|---|---|---|---|---|---|---|
| Title 1 | Film & Animation | 07 Sept. | 307 | 4.428 | 7 | 1419 | 2 | 17 |
| Title 2 | Music | 01 Nov. | 274 | 4.466 | 15 | 2992 | 4 | 32 |
| Title 3 | People & Blogs | 16 June | 46 | 3.428 | 98 | 20743 | 22 | 196 |
| Title 4 | News & Politics | 13 Aug. | 475 | 4.765 | 47 | 6117 | 24 | 95 |

automatically retrieved using the YouTube APIs.[2] Table 1 lists the title of the video and selected meta-data[3] for a few sample videos belonging to the input set.[4]

### 2.3 Video and User Properties

YouTube has two main objects: videos and users. We compute descriptive statistics of various video and user attributes belonging to the input set of videos (refer to Figure 2). Some insights that we draw from these descriptive statistics are:

1. We computed statistical measures for view counts (mean = 10640), favorites (mean = 17.05), comments (mean = 116.9), average rating (mean = 3.94) and raters (mean = 53.25).
2. The video length ranges from a minimum value of 21 seconds to a maximum value of 646 seconds. We notice that 25% of the videos have duration of less than 2 minutes (1st quartile = 124 seconds) and 50% of the videos have duration of less than 255 minutes (Median = 250 seconds). The data is right skewed as majority of the videos are less than 4.5 minutes (3rd quartile = 336.5). Our results was not different from Reid et al. – average length of videos was 6 minutes and 32 seconds [11].
3. The total numbers of favorites across all the videos are 1279. The mean value for number of favorites is 17.05. We notice that the total number of favorites for the top 5 videos is 518. This shows that there are certain videos which are very popular and are favorited by many users.
4. We found that 88% of the uploaders of videos are male and the average age is 25.4 (based on the information reported by users on their public profile). Our findings were aligned with previous studies by Adam et al. and Conway et al [1], [6].

---

[2] YouTube Data API http://code.google.com/apis/YouTube/

[3] The table displays values retrieved as of 11th May 2010.

[4] By the time we finished our analysis (June 23, 2010), we found that two of the videos were removed from YouTube because of violation of terms of service. This also confirms that the videos that we were studying are truly hate promoting videos.

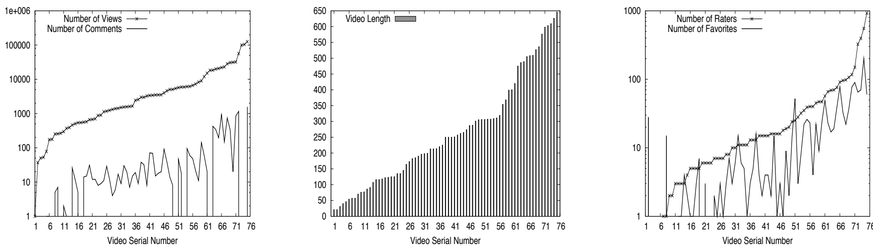


**Fig. 2.** Left: Presents number of views and comments for the videos in the data set; Middle: Presents video lengths of the videos in the data set; Right: Provides number of raters and number of favorites for the videos in the data set. The serial number of videos (x-axis) across the graph is not in the same order and is arranged with respect to monotonically increasing Y-axis value.

### 2.4 Linguistic Analysis of User Comments

We created a video-user matrix to identify users (from a set of 3037 unique users) who have commented on several videos belonging to the 75 extremist videos. Our hypothesis is that users who have actively commented (irrespective of the polarity or sentiment of the comment) on the extremist videos are potential candidates for further analysis. We identified top 12 most active users for the input set of videos and manually inspected their profile on YouTube to validate our hypothesis. We observed that 8 out of the 12 users are either active in promoting their ideologies or belong to the category of followers. We extract content-bearing words and phrases based on their presence across video comments. The extracted words and phrases consists of country and state names: India, Pakistan, Kashmir, religions: hindu, muslim, abusive phrases like: fuck you, shut up, son of a bitch and words like world, people and army. We performed an analysis of psychometric properties of users' comments using LIWC[5] (Linguistic Inquiry and Word Count) and topic discovery using LDA[6] (Latent Dirichlet Allocation) in order to gain a deeper understanding of textual messages. The output of LIWC reveal a high frequency of religious and swear words. The vocabulary size (number of unique words) of the user comments corpus was 2200 after eliminating all the stop words and symbols. We applied LDA to model documents (where all user comments for a video represents a document) as a mixture of topics where topics are distributions over words. We generated 3 topics from the comment corpus: Topic 1: miltary, dosti, succeed, jang, allaho, akabar; Topic 2: condom, viagara, penis, fuckpakigirl, hindoo, pundit; and Topic 3: mujahiddeen, hadith, jihad, saeed, kafir, destroy. Topic 2 contains terms which are more sexually abusive and derogatory, whereas Topic 3 contains words which are about terrorism and war.

[5] Linguistic Inquiry and Word Count (LIWC) http://www.liwc.net/
[6] Gensim Python Framework for Vector Space Modeling.

### 2.5 Social Network Analysis

**User Network (Friends) and Related Videos.** We created a set of unique userids who uploaded videos in the input set of 75 videos. A user can create a profile and invite other users to become friends. A particular user's friends list can be public or private depending on the users profile settings. If a user has not marked his or her friends list as private then the friends list can be viewed on the profile page of the respective user and can also be retrieved using the YouTube APIs. However, we notice that several users (33%) make their friend list as private and hence we were not able to retrieve friend list for such users. We created a social network graph only between the uploaders of videos belonging to our input set to test our hypothesis of the presence of a hidden virtual community and presence of influential or central users. Figure 3 (Left) is a social network graph which shows 60 unique users (derived from the input set) connected to each other using friends relationship. Only 25 of the 60 users have atleast one edge. It shows the presence of a hidden community of users having a shared interest and common agenda. Note that we are able to extract friend relationship only for those users who have marked their friend list as public and despite this restriction we observe a presence of a hidden community of users having a common interest. We compute statistical measures[7] such as betweenness centrality, closness centrality and degree to identify important and central nodes (leaders or influential users). Table 2 lists the top three userids (in decreasing order of rank) derived from computing statistical measures indicating their importance in the graph Figure 3 (Left).

Figure 3 (Right) presents a graph where each node represents a video in the input set of videos and each edge represents a related-video relationship. YouTube computes a list of related videos for each video based on the similarity of title, description, keywords and factors internal to YouTube. In the graph, two vertices are connected to each other if one video appears in the Top-25 related-video (based on the YouTube relevance ranking algorithm) list of the other video. Chatzopoulou et al. study related-video graph to understand the general characteristics and features of YouTube video [3]. We found various central videos fulfilling different purposes of hate-community for example Video ID: WV7 (maximum degree centrality) is a hate-propagating lecture by a so called "Dr." whereas Video ID: 0wm (maximum closeness centrality) shows brutalities against a section of people in India and thus incites anger. Video ID: 1rE (maximum betweeness centrality) supports terrorism and openly calls for a war.

**Multiple Relations between Users.** Figure 4 (Left) embodies in it all three relationships namely friends, subscription and video-shared. There is an edge from node A to node B if: A and B are *friends* and have mutually agreed to share content, A has *subscribed* to B and hence all B's updates are available to A or A has *favorited* or added a video in his playlist which has been uploaded by node B. In Figure 4 (Left), we observe that 40 out of 60 are connected through one of the three mentioned relationships. The layout in the figure has in the center the node

[7] Using JUNG: http://jung.sourceforge.net/

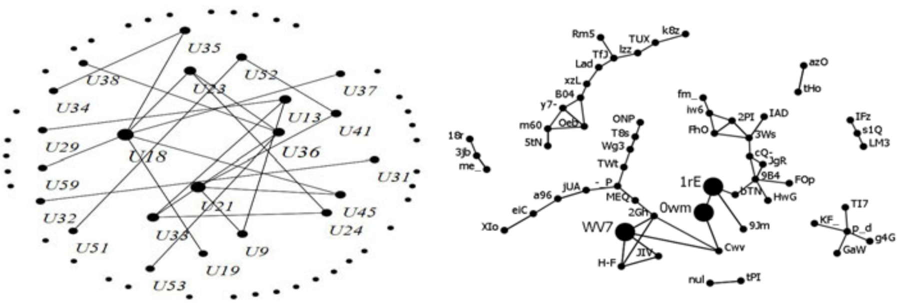

Fig. 3. Left: Social network graph (friends) between users of input set; Right: Presents the related video relationship. Only the first three characters of the video ID is shown in the figure and the central videos has been highlighted. We used SocNetV[8] and ORA[9] for creating these graphs.

**Table 2.** Top ranked users from social network graph in Figure 5 in terms of statistical measures indicating importance. We found some of the top 10 users are prominent in all measures. We see the user u18 in all three columns showing that he is most popular in the users that are studying.

| Rank | Betweenness Centrality | Closeness Centrality | Degree |
|---|---|---|---|
| 1 | u18 | u31 | u18 |
| 2 | u36 | u32 | u21 |
| 3 | u21 | u18 | u36 |

with highest betweeness centrality and it diminishes with radius. We found that user u36 stands out as a central leader. Centrality measures namely betweeness, closeness and inverse closeness statistically indicate that the topology of the network is core periphery (alpha=2).

**Ego-centric Network Graph Around Central Nodes.** Figure 4 (Right) presents an ego-centric network graph[10] (where the edge represents a bidirectional relationship of friends) for a user in the top 3 rank with respect to betweenness centrality and degree. Ego-centric is a graph which is centered on a particular vertex or node drawn to pay close attention to the relationship of a particular node (understand the view of network through the eyes of the "node.") It pays close attention to relationships of that node and the community structure around it. The graph in Figure 4 is an ego-centric graph where the depth is 2 and the maximum number of friends explored for a particular node is 50 (for illustration). Our hypothesis is that users who have high centrality in the socio-centric graph of the graph can be an entry point to further identify users having common interests and views (a belief in extremism and radicalization in our specific case

[8] http://socnetv.sourceforge.net/
[9] http://www.casos.cs.cmu.edu/projects/ora/
[10] Drawn using Vizster http://hci.stanford.edu/jheer/projects/vizster/

study). The basic premise is that a social network and ties of a user reflects the profile and interest of the user. Hence, community of persons which have high centrality in the socio-centric graph drawn from the uploaders of the input set of videos can reveal more persons with similar agenda and beliefs. We perform a manual inspection of the YouTube activity and profile of each node connected to the ego-center of Figure 4 (Right). Our analysis reveals that the ego-center is surrounded by people of the same kind or persons having common interests. We notice that the ego-center in the graph has 37 contacts (the communities or clusters are shaded) and a manual inspection of all the profiles reveals that 31 of the 37 contacts have YouTube activity which denotes hate and extremism. Amongst the remaining 6 contacts, 2 accounts were suspended (hence we cannot make a conclusion on these accounts but according to the YouTube policy an account is suspended if it violates community guidelines and terms of use).

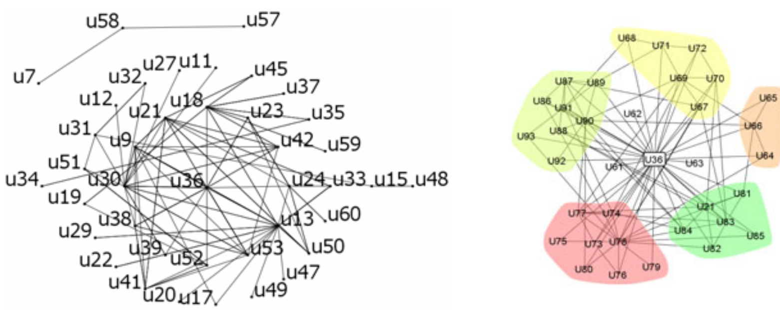


**Fig. 4.** Left: Social network graph (multiple relations) between users of input set. Right: Egocentric graph of an influential node (depth = 2, maximum friends = 50). The boundaries show communities that we found in the data.

**Subscription Relationship Network.** In contrast to friends relationship which is bidirectional, YouTube provides a unidirectional connection called as subscriptions. One user can subscribe to the videos uploaded by another user by subscribing to his channel. Unlike friends relationship, analysis of the network of users using subscription relationship is relatively unexplored in the literature. Biel et al. and Maia et al. recently studied the network of subscriptions and concluded that exploiting subscriptions relationship can offer additional insights into the patterns of users' behavior [2], [8]. Building on this, we hypothesize that additional hate-promoting users and contents can be discovered from the initial set of videos and uploaders by exploiting the subscription relationship. Our rationale is that if several users (who are already labeled as hate-promoting) subscribe to a particular user then it is highly likely that the subscribed user will share common interest. Similarly, if a user subscribes to many users in the input set of users then the subscriber is likely to have common interest with the users in the input set.

Figure 5 (Left) shows a selected portion of the subscription network for the users in the input set. The direction of arrow (since subscription relationship is a directed relationship) shows the flow of information which is opposite to the direction of subscription. The subscription network in Figure 5 (Left) reveals that some users are a major source of content provider for the users in the input set. This observation can be used to uncover additional users and community resulting in bootstrapping from the entry point. We validate our conjecture by performing a manual inspection of the profile and activity of the discovered users.

Maia et al. studied subscription relationship and interaction patterns between users in YouTube to characterize and identify user behavior [8]. We draw from Maia et al. idea of categorizing user behavior based on subscription activity and identify 10 users who have maximum number of subscriptions. These 10 users can be classified as hate-promoting content seekers. A node with high out-degree or large number of subscribers indicates a content producer. We identify 10 users who have been subscribed by the initial set of users and act as information hubs for the hate community. We identified four users who are common in both the list which can be categorized as nodes who are playing active role in both dissemination and consumption of hate content.

**Favorite/Playlist Relationship Network.** Figure 5 (Right) presents connections between users in the input dataset and the videos favorited/playlisted (added to their playlist) by them. Both of these action are like-video actions. Our hypothesis is that if several users (who have been tagged as hate-promoting) favorite/playlist a particular video then the likelihood of the liked video being hate-promoting is high. The rest set of arrows in the left in Figure represents edges connecting a user and a video through a favorite/playlist relationship. The second set of arrows on the right connects a video and its uploader. We found three of the top six (max. in-degree) videos had a clear hate-promoting agenda. The rest three videos were not hate-spreading (through the video content) but were sensitive as they received several hateful comments. A careful analysis of the uploaders profile reveals that four uploaders amongst the six belonged to the hate-promotion category. None of these four uploaders as well as the six

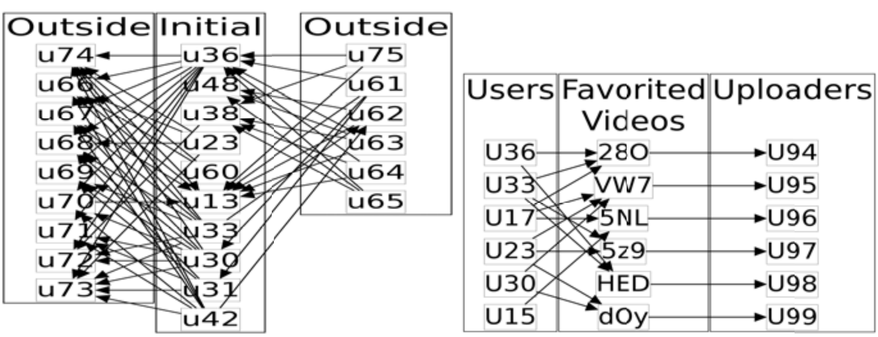


**Fig. 5.** Left: Shows a selected portion of the subscription network of users. One in the middle are from our data set, ones in the left and right of this figure are outside our data set; Right: Presents a connection between users in the input dataset and the videos favorited by them.

videos were present in the initial seed list of users and hence we were able to augment the list of the hate-promoters. This supports our hypothesis that favorite/playlist relationship can be exploited to bootstrap the initial set of videos and users.

**Iterative Expansion of Seed List of Users and Community.** Table 3 present empirical results to test the proposed hypothesis by combining three relations (friends, subscriptions and favorites) as a single relation and then expanding the user network graph from a seed list of users to identify additional like-minded users and community. The system extends the seed list of users in each iteration based on five centrality measures: in-degree, hub, information, out-degree and betweeness. We measure the precision of the system by manually validating user's profiles. The system was able to add 98 (true positive) users with an average precision of 88% in two iterations.

**Table 3.** Empirical results of the bootstrapping process (addition of 98 users from the initial seed list of users within 2 iterations). Abbreviations - Iter: Iteration, TP: True Positive, FP: False Positive, CD: Can't Determine, Prec: Precision, NU: New Users. Top-K represents the number videos that we took for the analysis after ranking them.

| Iter | Seed | Nodes | Links | Centrality | Top-K | TP | FP | CD | Prec | NU |
|---|---|---|---|---|---|---|---|---|---|---|
| 1 | 60 | 1628 | 5649 | Indegree | 50 | 46 | 4 | 0 | 0.92 | 19 |
| | | | | Hub | 50 | 48 | 2 | 0 | 0.96 | 23 |
| | | | | Information | 50 | 48 | 2 | 0 | 0.96 | 16 |
| | | | | Outdegree | 50 | 48 | 2 | 0 | 0.96 | 16 |
| | | | | Betweeness | 50 | 40 | 7 | 3 | 0.85 | 18 |
| 2 | 106 | 5240 | 30481 | Indegree | 100 | 88 | 11 | 1 | 0.88 | 23 |
| | | | | Hub | 100 | 85 | 13 | 2 | 0.86 | 36 |
| | | | | Information | 100 | 88 | 10 | 2 | 0.89 | 31 |
| | | | | Outdegree | 100 | 92 | 7 | 1 | 0.92 | 25 |
| | | | | Betweeness | 100 | 83 | 14 | 3 | 0.85 | 12 |
| Avgerage Precision : 0.88 | | | | | New Users (TP) Added: 98 | | | | | |

## 3 Discussion

We present an approach (based on exploiting relations between users) to retrieve hate and extremist videos, users and communities from YouTube. The proposed system was able to bootstrap from 60 (seed-list) to 158 (true positive) users in two iterations. The system was able to search 98 users automatically with a precision of 88%. The proposed approach can discover central, and influential users and videos as well as hidden communities using social network analysis techniques (using friends, subscriptions, favorites and related videos). The output shows that the proposed approach can potentially help a security analyst find what he is looking for (i.e., able to assist in solving his information need) and producing an output in a form (reports, network graphs) which is more insightful and actionable than just a flat list of videos and users.

## Acknowledgments

This work was supported by a DIT (Department of Information Technology), Government of India grant. The authors would like to thank all members of PreCog research group at IIIT-Delhi.

## References

1. Bermingham, A., Conway, M., McInerney, L., O'Hare, N., Smeaton, A.F.: Combining Social Network Analysis and Sentiment Analysis to Explore the Potential for Online Radicalisation. In: IEEE International Conference on Advances in Social Network Analysis and Mining, Washington, DC, USA, pp. 231–236 (2009)
2. Biel, J.: Please, Subscribe to Me! Analysing the Structure and Dynamics of the YouTube Network
3. Chatzopoulou, G., Sheng, C., Faloutsos, M.: A First Step Towards Understanding Popularity in YouTube. In: Second International Workshop on Network Scienece for Communication Networks (NetSciCom), San Diego, USA (2010)
4. Chau, M., Xu, J.: Mining Communities and their Relationships in Blogs: A Study of Online Hate Groups. Int. J. Hum.-Comput. Stud. 65(1), 57–70 (2007)
5. Chen, H., Chung, W., Qin, J., Reid, E., Sageman, M.: Uncovering the Dark Web: A Case Study of Jihad on the Web. Journal of American Society for Information Science and Technology 59(8), 1347–1359 (2008)
6. Conway, M., Mcinerney, L.: Jihadi video and Auto-radicalisation: Evidence from an Exploratory YouTube Study. In: Ortiz-Arroyo, D., Larsen, H.L., Zeng, D.D., Hicks, D., Wagner, G. (eds.) EuroIsI 2008. LNCS, vol. 5376, pp. 108–118. Springer, Heidelberg (2008)
7. Fu, T., Huang, C.-N., Chen, H.: Identification of Extremist Videos in Online Video Sharing Sites. In: ISI 2009: Proceedings of the 2009 IEEE International Conference on Intelligence and Security Informatics, Piscataway, NJ, USA, pp. 179–181. IEEE Press, Los Alamitos (2009)
8. Maia, M., Almeida, J., Almeida, V.: Identifying User Behavior in Online Social Networks. In: SocialNets 2008: Proceedings of the 1st Workshop on Social Network Systems, pp. 1–6. ACM, New York (2008)
9. Mislove, A., Marcon, M., Gummadi, K.P., Druschel, P., Bhattacharjee, B.: Measurement and Analysis of Online Social Networks. In: IMC 2007: Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement, pp. 29–42. ACM Press, New York (2007)
10. Paolillo, J.C.: Structure and Network in the YouTube Core. In: HICSS 2008: Proceedings of the Proceedings of the 41st Annual Hawaii International Conference on System Sciences, p. 156. IEEE Computer Society Press, Los Alamitos (2008)
11. Reid, E.: Analysis of Jihadi Extremist Groups Videos. Forensic Science Communications 11(3) (2009)
12. Santos, R.L.T., Rocha, B.P.S., Rezende, C.G., Loureiro, A.A.F.: Characterizing the YouTube video-sharing community