

Provided by the author(s) and University College Dublin Library in accordance with publisher policies. Please cite the published version when available.

Title	Down the (White) Rabbit Hole: The Extreme Right and Online Recommender Systems
Author(s)	O'Callaghan, Derek; Greene, Derek; Conway, Maura; Carthy, Joe; Cunningham, Pádraig
Publication date	2015
Publication information	Social Science Computer Review, 33 (4): 459-478
Publisher	Sage Publications
Item record/more information	http://hdl.handle.net/10197/8439
Publisher's version (DOI)	http://dx.doi.org/10.1177/0894439314555329

Downloaded 2018-04-21T17:53:52Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa) 

Some rights reserved. For more information, please see the item record link above.



Down the (White) Rabbit Hole: The Extreme Right and Online Recommender Systems

Derek O'Callaghan (UCD)¹ Derek Green (UCD), Maura Conway (DCU)², Joe Carthy (UCD), and Pádraig Cunningham (UCD)

In addition to hosting user-generated video content, YouTube provides recommendation services, where sets of related and recommended videos are presented to users, based on factors such as co-visitation count and prior viewing history. This article is specifically concerned with extreme right (ER) video content, portions of which contravene hate laws and are thus illegal in certain countries, which are recommended by YouTube to some users. We develop a categorization of this content based on various schema found in a selection of academic literature on the ER, which is then used to demonstrate the political articulations of YouTube's recommender system, particularly the narrowing of the range of content to which users are exposed and the potential impacts of this. For this purpose, we use two data sets of English and German language ER YouTube channels, along with channels suggested by YouTube's related video service. A process is observable whereby users accessing an ER YouTube video are likely to be recommended further ER content, leading to immersion in an ideological bubble in just a few short clicks. The evidence presented in this article supports a shift of the almost exclusive focus on users as content creators and protagonists in extremist cyberspaces to also consider online platform providers as important actors in these same spaces.

Keywords: extreme right, categorization, recommender systems, topic modeling, YouTube

Introduction

For youth, the Internet is often their first port of call for background and information on topics with which they are unfamiliar or, indeed, for discussion and networking around topics with which they are (Agosto & Hughes-Hassell, 2006a, 2006b; Ito et al., 2010; see also Livingstone & Haddon, 2009). Political extremists are aware of this trend and seek to exploit it through the use of not just dedicated websites but also by pushing out their content across the whole of the Internet, including via social media platforms such as Facebook, Twitter, and YouTube (Conway, 2012; Conway & McInerney, 2008; Europol, 2011, pp. 11–12; U.K. Home Office, 2011, p. 35). In this way, they aim to reach a much wider audience than they previously had access to, which may be a factor in explaining the high numbers of European children and teens who report having accessed hate content via the Internet: a major study found that some 12% of European 11- to 16-year-olds reported seeing online hate in the year prior to interview, rising to one in five 15- to 16-year-olds (Livingstone, O'lafsson, O'Neill, & Donoso, 2012, p. 11). YouTube's status as the most popular video sharing platform means that it is especially useful to political extremists. In addition to hosting user-generated video content, YouTube provides recommendation services, where sets of related and recommended videos are presented to users, based on factors such as co-visitation count and prior viewing history (Davidson et al., 2010).

¹ University College Dublin, Dublin, Ireland

² Dublin City University, Dublin, Ireland

YouTube, and other social media sites, expends a lot of effort on making “strategic claims for what they do and do not do, and how their place in the information landscape should be understood” (Gillespie, 2010, p. 347). However, online platform providers are facing increasing scrutiny from users, advertisers, activists, policy makers, and other key constituencies about their civic responsibilities including, in particular, along a key fault line between intervening in online content delivery versus remaining neutral. YouTube have been at pains “to position themselves as just hosting—empowering all by choosing none” (Gillespie, 2010, p. 357; see also Fiore-Silfvast, 2012, p. 1982). They insist that:

YouTube encourages free speech and defends everyone’s right to express unpopular points of view. We believe that YouTube is a richer and more relevant platform for users precisely because it hosts a diverse range of views, and rather than stifle debate we allow our users to view all acceptable content and make up their own minds’ [author italics]. (YouTube Team, 2008)

In fact, as critical media scholars point out, social media platforms do not merely transmit content, but filter it on the basis of claiming to augment it, thereby making the content more relevant to its potential consumers (and the platforms more attractive to advertisers; see, e.g., Bucher, 2012; Langlois&Elmer, 2013). Tania Bucher has shown that “APIs have ‘politics’” (2013, p. 2); in this article, we show that recommender systems also have politics. This means that they too can be seen as having “powerful consequences for the social activities that happen with them, and in the worlds imagined by them” (Gillespie, 2003, as quoted in Bucher 2013, p. 2). In particular, the YouTube recommender system ensures that, contrary to YouTube’s own assertions, users are explicitly not exposed to “all acceptable content.”

Whereas previous studies in this area have highlighted the online ideological bubbles or echo chambers resulting from choices made by content consumers (Pariser, 2011; Sunstein, 2001), this article is concerned with content recommenders, in terms of the video and channel links suggested by YouTube. The focus herein is extreme right (ER) content, such as music and other associated propaganda, portions of which are deemed hate content and thus illegal in certain countries, which is made freely available via YouTube, often for long periods of time (Bell, 2013). Our analysis is concerned with two aspects of the latter content; first, what is being posted (i.e., the nature of YouTube video and channel content) in order to categorize it for the purposes of showing, second, how automated social media targeting or “recommendation” can have the undesirable consequence of a user being excluded from information that is not aligned with their existing perspective, potentially leading to immersion within an ideological bubble. We show that this is, in fact, an almost certain outcome of users accessing an ER channel on YouTube. Using data on English and German language ER content, we show that YouTube users are very likely to be recommended further ER content within the same category or related ER content from a different category but are unlikely to be presented with non-ER content. This suggests that YouTube’s recommender algorithms are not neutral in their effects but have political articulations (Slack & Wise, 2007, Chap. 11).




Figure 1. Overview of the process used to investigate the extreme right content found on YouTube

The article is divided into four sections. The first section provides description of prior work on YouTube categorization and recommendation, along with a review of research concerned with the categorization of online ER content. The retrieval of YouTube data based on links originating from ER Twitter accounts is the subject of the second section, while in the third section, we describe the methodology used for related channel ranking, topic identification, and channel categorization. An overview of this entire process can be found in Figure 1. Our investigation into recommender politics and its potential impacts can be found in the fourth section, which is focused on two data sets consisting primarily of English and German language channels, respectively. In the Discussion section, we emphasize the hidden politics of online recommender systems, particularly the way in which the immersion of some users in YouTube's ER spaces is a coproduction between the content generated by users and the affordances of YouTube's recommender system, and the potential implications of and suggested responses to this. The Conclusion synthesizes our research and findings and makes a suggestion for future research.

Related Research

Video Recommendation and Categorization

Video recommendation on YouTube has been the focus of a number of studies. For example, Baluja et al. (2008) suggested that standard approaches used in text domains were not easily applicable due to the difficulty of reliable video labeling. They proposed a graph-based approach that utilized the viewing patterns of YouTube users, which did not rely on the analysis of the underlying videos. The recommendation system in use at YouTube at the time was discussed by Davidson et al. (2010) where sets of personalized videos were generated with a combination of prior user activity (videos watched, favoured, liked) and the traversal of a co-visitation graph. In this process, recommendation diversity was obtained by means of a limited transitive closure over the generated related video graph. Zhou, Khemmarat, and Gao (2010) performed a measurement study on YouTube videos to determine the sources responsible for video views and found that related video recommendation was the main source outside of the search function for the majority of videos. They also found that the click-through rate to related videos was high, where the position of a video in a related video list played a critical role. A similar finding was made by Figueiredo, Benevenuto, and Almeida (2011) where they demonstrated the importance of key mechanisms such as related videos in the attraction of users to videos.

Turning to the task of YouTube video categorization, Filippova and Hall (2011) presented a text based method that relied upon metadata such as video title, description, tags and comments, in conjunction with a predefined set of 75 categories. Using a bag-of-words model, they found that all of the text sources contributed to successful category prediction. More recently, a framework for the categorization of video channels was proposed by Simonet (2013), involving the use of semantic entities identified within the corresponding video and channel profile text metadata. Following the judgment that existing taxonomies were not well suited to this particular problem, a new category taxonomy was developed for YouTube content. Roy, Mei, Zeng, and Li (2012) investigated both video recommendation and categorization in tandem, where videos were categorized according to the topics built from Twitter activity, leading to the enrichment of related video recommendation. Video text metadata were used for this process, and the topics were based on the categories proposed by Filippova and Hall (2011) in addition to the standard YouTube categories at the time. Separately, they also analyzed diversity among related videos, where they found that there was a 25% probability on average of a related video (up to a depth of 2) being in the same category.

ER Categorization

In the various studies that have analyzed online ER activity, certain differences can be observed among researchers in relation to their categorization of this activity and associated organizations (Blee & Creasap, 2010; see also Table 1). Burris, Smith, and Strahm (2000) proposed a set of eight primarily U.S.-centric categories in their analysis of a White supremacist website network; Holocaust Revisionists, Christian

Identity Theology, Neo-Nazis, White Supremacists, Foreign (non-US) Nationalists, Racist Skinheads, Music, and Books/Merchandise. A similar schema was used by Gerstenfeld, Grant, and Chiang (2003), which also included Ku Klux Klan and Militia categories. They also discussed the difficulty involved in the categorization of certain subgroups, where a general category (Other) was applied in such cases. These categories were adapted in separate studies of Italian and German ER groups, where new additions included Political Parties and Conspiracy Theorists, while others such as Music and Skinheads were merged into a Young category (Caiani & Wagemann, 2009; Tateo, 2005). Rather than focusing on ideological factors, Goodwin and Ramalingam (2012) proposed four organizational types found within the European ER milieu; political parties, grassroots social movements, independent smaller groups, and individual “lone wolves.” Other notable categories include the Autonomous Nationalists identified within Germany in recent years. These groups focus specifically on attracting a younger audience, where social media is often a critical component in this process (Baldauf, Groß, Rafael, and Wolf, 2011).

The popularity of YouTube has led to its usage by ER groups for purposes of content dissemination. Its related video recommendation service provides a motivation for the current work to analyze the extent to which a viewer may be exposed to such content and thereby highlighting “Whether these interventions are strategic or incidental, harmful or benign, they are deliberate choices that end up shaping the contours of public discourse online” (Gillespie, 2010, p. 358). Separately, disagreements over the categorization of online ER activity suggest that a specific set of categories may be required for the analysis of this domain.

Data

In our previous work, we investigated the potential for Twitter to act as a gateway to communities within the wider online network of the ER (O’Callaghan, Greene, Conway, Carthy, & Cunningham, 2013b). Two data sets associated with ER English language and German language Twitter accounts were generated by retrieving profile data over an extended period of time. We gathered all tweeted links to external websites and used these to construct an extended network representation. In the current work, we are solely interested in tweeted YouTube URLs. Data for these Twitter accounts were retrieved between June 2012 and May 2013, as limited by the Twitter REST application programming interface (API) restrictions effective at the time. YouTube URLs found in tweets were analyzed to determine a set of channel (synonymous with uploaders or accounts; see Simonet, 2013) identifiers that were directly (channel profile page URL) or indirectly (URL of video uploaded by channel) tweeted. All identified channels were included, regardless of the number of tweets in which they featured. Throughout this work, we refer to these as seed channels; 26,460 and 3,046 were identified for the English language and German language data sets, respectively.

Our first step in this investigation is the categorization of YouTube channels. To do this, we use text metadata associated with the videos uploaded by a particular channel—namely, their titles, descriptions, and associated key words. Although user comments have been employed in other work (Filippova & Hall, 2011), they were excluded here. This decision followed an initial manual analysis of a sample of tweeted videos, which found that comments were often not present or had been explicitly disabled by the uploader. We also excluded the YouTube “category” field, as it was considered too broad to be useful in the ER domain. Using the YouTube Data API (see <https://developers.google.com/youtube/>), we initially retrieved the available text metadata for up to 1,000 of the videos uploaded by each seed channel, where the API returns videos in reverse chronological order according to their upload time. In cases where seed channels and their videos were no longer available (e.g., the channel had been suspended or deleted since appearing in a tweet), these channels were simply ignored.

Table 1. Categories of Extreme Right YouTube Content, Based on Common Categorizations Found in Academic

Category	Description	Source
Anti-Islam	Can include political parties (e.g., Dutch PVV) or groups such as the English Defence League (EDL), which often describe themselves as “counter-Jihad”	Baldauf, Groß, Rafael, and Wolf (2011); Goodwin (2013)
Anti-Semitic	All types of anti-Semitism, regardless of association (existing literature tends to discuss this in relation to other categories)	Burris, Smith, and Strahm (2000); Tateo (2005)
Conspiracy Theory	Themes include New World Order (NWO), Illuminati, etc. Not exclusively ER, but often related to Patriot in this context	Baldauf et al. (2011); Southern Poverty Law Center (SPLC) 2013; Tateo (2005)
Music	Includes any ER music such as Oi!, Rock Against Communism (RAC), etc.	Baldauf et al. (2011); Burris et al. (2000); Tateo (2005)
Neo-Nazi	Nazi references, such as to Hitler, WWII, SS, etc.	Baldauf et al. (2011); Burris et al. (2000); Gerstenfeld, Grant, and Chiang (2003)
Patriot	U.S.-centric, including groups such as “Birthers,” militia, antigovernment, anti-immigration, opposition to financial system. Some of these themes are not exclusive to ER	SPLC 2013
Political Party	Primarily European parties such as the BNP, FPO, Jobbik, NPD, PVV, Swedish Democrats, UKIP, etc. Many of these parties are also categorized as Populist	Baldauf et al. (2011); Caiani and Wagemann (2009); Goodwin and Ramalingam (2012)
Populist	Broader category that includes various themes such as anti-EU, antiestablishment, antistate/government, anti-immigration (as with Patriot, some of these are not exclusive to ER). Although some disagreement about this category exists (Marlie`re, 2013), we have used it, as it has proved convenient for categorizing certain groups that span multiple themes	Bartlett, Birdwell, and Littler (2011); Mudde, (2007)
Revisionist	References to Holocaust/WWII denial. Closely associated with Neo-Nazi	Burris et al. (2000); Gerstenfeld et al. (2003); Tateo (2005)
Street Movement	Groups such as the EDL, Autonome Nationalisten, Spreelichter, Anti-Antifa, etc.	Baldauf et al. (2011); Goodwin and Ramalingam (2012)
White Nationalist	References to white nationalism and supremacism, also used to characterize political parties such as the BNP or Jobbik	Burris et al. (2000); Gerstenfeld et al. (2003); Tateo (2005)

Note. PVV = Partij voor de Vrijheid; EDL = English Defence League; NOW = NewWorld Order; ER=extreme right; WWII=World War II; SS = Schutzstaffel; BNP = British National Party; FPO = Freedom Party of Austria; NPD = National Democratic Party of Germany; UKIP = United Kingdom Independence Party; EU = European Union. We have employed multiple sources due to the fact that no definitive set of categories is agreed upon in this domain. See also the “Extreme Right Categorization” and “Topic and Channel Categorization” sections.

To address the variance in the number of uploaded videos per channel, and to reduce the volume of subsequent data retrieval, we randomly sampled up to 50 videos for each seed. For each video in this sample, metadata values were retrieved for the top 10 related videos returned by the API. Using the default parameter settings, these videos appear to be returned in order of “relevance,” as defined internally by YouTube, similar to the default behavior of video search results feeds described in the API documentation (June 2013). We refer to the corresponding uploaders as related channels; 1,451,189 and 195,146 were

identified for the English and German data sets, respectively. As before, we then retrieved the available text metadata for up to 1,000 videos uploaded by each unseen related channel, from which a random sample of up to 50 videos was generated. Separately, YouTube automatically annotates, where possible, uploaded videos using topics found in Freebase (see <http://www.freebase.com/>), a crowd sourced collection of structured data on a large multiplicity of topics (for details of the YouTube annotation process, using Freebase, see http://youtu.be/wf_77z1H-vQ). We therefore also retrieved all available Freebase topic assignments for the sampled videos.

For both seed and related channels, the corresponding uploaded video sample was used for categorization; this is described subsequently. It is worth noting here that the videos in question may have been uploaded at any time prior to retrieval, where these times are not necessarily restricted to the period of either Twitter or YouTube data retrieval.

Method

Having retrieved the channel and video data from YouTube (part (a) of the process overview diagram found in Figure 1), the next steps were concerned with the processing of these data prior to the investigation of YouTube’s recommender system. This section corresponds to part (b) of the diagram in Figure 1. First, a ranking of related channels was generated for each seed channel, using the singular value decomposition (SVD) rank aggregation method proposed by Greene and Cunningham (2013). Next, the channels were represented as “documents” based on an aggregation of their uploaded video text metadata, within which latent topics were identified with Nonnegative Matrix Factorization (NMF; Lee & Seung, 1999). The channels were then categorized using these topics in combination with a proposed categorization based on various schema found in a selection of academic literature on the ER, where a simple example of the process that maps channels to categories is shown in Figure 2.




Figure 2 Topic and channel categorization process

Related Channel Ranking

It was first necessary to generate a single related channel ranking for each seed channel, according to the multiple related rankings returned by the YouTube API for the sample of videos uploaded by the seed. The SVD rank aggregation method proposed by Greene and Cunningham (2013) was used to combine the rankings for each video uploaded by a particular seed into a single ranked set, from which the top ranked related channels were then selected. It was decided to restrict the focus to the top 10, given the impact of related video position on click-through rate (Zhou, Khemmarat, & Gao, 2010). These are the channels used for our detailed investigation into the political articulations of YouTube's recommender algorithm given subsequently.

Topic Identification

For the purpose of channel categorization, latent topics associated with the channels in both data sets were initially identified based on their uploaded videos. Following the approach of Hannon, Bennett, and Smyth (2010), a "profile document" was generated for each seed and related channel, consisting of an aggregation of the text metadata from their corresponding uploaded video sample, from which a tokenized vector representation was produced. Topic modeling is concerned with the discovery of latent semantic structure or topics within a text corpus, which can be derived from co-occurrences of words in documents (Steyvers & Griffiths, 2006). Popular methods include probabilistic models such as latent Dirichlet allocation (LDA; Blei, Ng, & Jordan, 2003) or matrix decomposition techniques such as NMF (Lee & Seung, 1999). Both LDA and NMF-based methods were initially evaluated with the seed channel document representations described. However, NMF was found to produce the most readily interpretable results, which appeared to be due to the tendency of LDA to discover topics that overgeneralized (Chemudugunta, Smyth, & Steyvers, 2006; O'Callaghan, Greene, Conway, Carthy, & Cunningham, 2013a). There was an awareness of the presence of smaller groups of channels associated with multiple languages in both data sets, and it was decided to opt for specificity rather than generality. The process was undertaken in two stages. First, topics for the seeds were generated by applying NMF to the seed channel documents, resulting in a set of basis vectors consisting of both ER and non-ER topics for both data sets. The second stage involved producing assignments to these seed topics for the related channel documents.

Topic and Channel Categorization

Some prior work discussed earlier has proposed generic categories for use with YouTube videos and channels (Filippova & Hall, 2011; Roy et al., 2012). However, as these studies focused on the categorization of mainstream videos, they were not sufficient for the present analysis where categories specifically associated with the ER were required. We have also discussed prior work that characterized online ER activity using a number of proposed categories but, as indicated, no definitive set of categories is agreed upon in this domain (Blee & Creasap, 2010). We therefore propose a categorization based on various schema found in a selection of academic literature on the ER, where this category selection is particularly suited to the ER videos and channels we have found on YouTube. Some categories are clearly delineated while others are less distinct, reflecting the complicated ideological makeup and thus fragmented nature of groups and subgroups within the ER (Gerstenfeld, Grant, & Chiang, 2003). In such cases, we have proposed categories that are as specific as possible while also accommodating a number of disparate themes and groups. Details of the categories employed can be found in Table 1.

As both data sets contained various non-ER channels, we also created a corresponding set of non-ER categories consisting of a selection of the general YouTube categories as of June 2013, in addition to other categories that we deemed appropriate following an inspection of these channels and associated topics. These non-ER categories were Entertainment, Gaming, Military, Music, News & Current Affairs, Politics, Religion, Science & Education, Sport, and Television. Having produced a set of T topics for a data set, we then proceeded to categorize them. For each topic, we manually inspected the high-ranking topic terms, in

addition to profiles and uploaded videos for a selection of seed channels most closely associated with the topic. Multiple categories were assigned to topics where necessary, as using a single category per topic would have been too restrictive while also not reflecting the often multifaceted nature of most topics that were identified. In many cases, categories for topics were clearly identifiable, with a separation between ER and non-ER categories. For example, an *English Defence League* (EDL) topic was categorized as Anti-Islam and Street Movement, while a topic having high-ranking terms such as “guitar” and “band” was categorized as Music. For certain topics, this separation was more ambiguous, where the channels associated with a particular topic consisted of a mixture of both ER and non-ER channels. A combination of both ER and non-ER categories was assigned in such cases.

The set of categorized topics was then used to label both seed and related channels, based on the channel topic assignment weights, as determined by NMF, exceeding a threshold w . This supports the potential assignment of multiple categories to a single channel. We found that a certain number of channels, both seed and related, had a flat profile with relatively low weights ($<w$) for all topics; these could be considered as *gray sheep*, a term used in the recommender systems literature to describe users whose opinions do not consistently agree or disagree with any groups (Claypool et al., 1999). Further analysis found this to be due to factors such as the original documents being short or containing few unique terms. As we were unable to reliably categorize such channels, they were excluded from all subsequent analysis. Separately, although the NMF process identified topics with high-ranking discriminating terms in languages other than that of the data set, a small number of topics with less discriminating general language terms were also found. As it would have been difficult to distinguish between ER and non-ER channels closely associated with these topics, these were also excluded. A simple example illustrating the complete process of mapping of channels to categories is shown in Figure 2. In order to confirm the reliability of the above NMF topic categorizations we employed the automatic Freebase topic annotation of videos by YouTube (hereafter referred to as F-topics to distinguish them from topics discovered by NMF). A set of member seed channels was determined for each NMF topic, where a seed channel was considered a member of each topic used in its categorization (topic-channel assignment weight $\geq w$). Next, an “F-topic document” vector was created for each seed channel, by aggregating the English-language labels for all F-topics assigned to their respective uploaded videos. Finally, for each NMF topic, the subset of vectors for the member seed channels was used to calculate a mean vector D , with the F-topic ranking consisting of the top 10 F-topic identifiers in D , while the mean pairwise Cosine similarity between the member vectors was also calculated. These F-topic rankings were then compared with the corresponding NMF topic categorizations.

The ER and the Political Articulations of YouTube’s Recommender System

In this section, we discuss the political articulations of the English and German language YouTube data sets. The experimental steps taken were as follows:

1. Generation of an aggregated ranking of related channels for each seed channel.
2. Generation of channel document vectors and identification of topics using NMF.
3. Categorization of the identified topics according to the set defined in Table 1.
4. Categorization of the channels based on their topic association weights.
5. Investigation of whether information was excluded from viewing by users if it was not aligned with their existing—in this case, ER—perspective, potentially leading to immersion within an ideological bubble.

For Step 5 above, we define an ER ideological bubble in terms of the extent to which the related channels for a particular ER seed channel also feature ER content. It has been shown that the position of a video in a related video list plays a critical role in the click-through rate (Zhou et al., 2010). Therefore, we investigated the existence of an ideological bubble using the top k ranked related channels, with increasing values of $k \in [1,3,5,7,10]$, as follows. For each ER seed channel:

1. The top k ranked related channels were selected, filtering any excluded channels as discussed previously. Seed channels with no remaining related channels following filtering were not considered for rank k .

2. The total proportion of each category assigned to the $\leq k$ related channels was calculated.

Then, for each ER category:

1. All seed channels to which the category has been assigned were selected.
2. The mean proportion of each related category associated with these seed channels was calculated.

We considered an ideological bubble to exist for a particular ER category when its highest ranking related categories, in terms of their mean proportions, were also ER categories.

English Language Categories

From the total number of channels in the English language data set, we generated 24,611 seed and 1,376,924 related channel documents, using a corresponding seed-based vocabulary of 39,492 terms, and topics were identified by applying NMF to the seed documents.

Table 2. Highest Ranked English-Language ER NMF Topic Categorisation and their Corresponding Freebase Topics, Based on YouTube Annotations of the Sampled Videos.

#Members	Mean Similarity	Category	Top 10 Freebase topics
88	0.33	Populist	Anonymous, Occupy movement, Occupy Wall Street, Ron Paul, Protest, Illuminati, Scientology, Documentary, New World Order, Barack Obama
257	0.30	Anti-Islam, Street Movements	English Defence League, Tommy Robinson, Unite Against Fascism, Muslim, Islam, Luton, British National Party, Leicester, Sharia, Combat
416	0.23	Music	Heavy metal, Black metal, Concert, Death metal, Thrash metal, Rock music, National Socialist black metal, Doom metal, Folk metal, Album
155	0.21	White Nationalist, Political Party, Populist	British National Party, Nick Griffin, Unite Against Fascism, English Defence League, Muslim, Islam, Racism, Member of the European Parliament, United Kingdom, Andrew Brons
55	0.18	Anti-Islam	George Galloway, Islam, Muslim, Israel, Christopher Hitchens, Hamas, Iran, English Defence League, Documentary, Yvonne Ridley
357	0.17	Anti-Semitic	Israel, Palestinian people, Gaza, Hamas, Israel Defense Forces, Zionism, Jewish people, Palestine, Islam, Gaza Strip
798	0.17	Anti-Islam	Islam, Muslim, Muhammad, Quran, Sharia, Israel, English Defence League, Jihad, Allah, Jesus Christ
220	0.11	Music	Oil, Punk rock, Skinhead, Concert, Rock music, Rock Against Communism, Alternative rock, Heavy metal, Hardcore punk, Theme music
952	0.11	Patriot	Barack Obama, Mitt Romney, Ron Paul, Alex Jones, Politics, Interview, Republican Party, Israel, Glenn Beck, Sarah Palin
155	0.11	Populist, Political Party	U.K. Independence Party, European Union, Nigel Farage, Member of the European Parliament, Europe, European Parliament, Muslim, Islam, Jewish people, Immigration

Note. These NMF topics contain ≥ 10 member seed channels and are ranked using mean pairwise Cosine similarity of the corresponding F-topic vectors. As discussed in Table 1, some of these issues are not exclusive to the extreme right, for example, the Occupy movement (considered populist in an ER context for the purpose of the current work), or various subgenres of metal music appearing with National Socialist Black Metal.

To determine the number of topics T , we experimented with values of T in $[10,100]$ to produce topics that were as specific as possible, given prior knowledge of the presence of smaller groups of channels associated with multiple languages within both data sets. This led to the selection of $T = 80$, as larger values resulted in topic splits rather than the emergence of unseen topics. Of these, 27 ER topics (33.75%), 39 non-ER topics (48.75%), 8 topics that were a combination of ER and non-ER categories (10%), and 6 topics based on general terms of a separate language (7.5%) were found. These 80 topics were then categorized according to the set defined in Table 1, which permitted the subsequent categorization of the seed and related channels using their corresponding topic assignment weights. At this point, we excluded 8,225 (33.42%) seed and 482,226 (35.66%) related gray sheep channels that could not be categorized. Related channels ranked at $k > 10$ were also excluded, and all non-ER seed channels were removed from the candidate seed set. The remaining 6,573 ER seed and their 22,980 related channels were used to calculate the mean related category proportions for each ER category. The reliability of the categorization process is confirmed by Table 2, which contains a comparison of ER NMF topic categorizations with the corresponding highest ranked Freebase topic assignments of the member seed channel videos (the mean percentage of sampled videos having annotated topics was 92% per channel, with $\sigma = 14\%$). As can be seen in Table 2, there is a high level of consistency between the two.




Figure 3 English mean category proportions of top k ranked related channels ($k \in [1,3,5,7,10]$), for selected seed ER categories. ER = extreme right; NER = non-extreme right.

Figure 3 contains plots for three ER categories that were selected for more detailed analysis, where ER and non-ER categories have been prefixed with ER- and NER, respectively. To assist visual interpretation, any weakly related categories whose mean proportion was $<.06$ for a particular k ranking have been omitted. From inspecting these plots, two initial observations can be made: (1) the seed category is the dominant related category for all values of k and (2) although related category diversity increases at lower k rankings with the introduction of certain non-ER categories, ER categories consistently have the strongest presence. In the case of Anti-Islam seed channels, it appears that the top ranked related channels ($k = 1$) are mostly affiliated with various groups from the United Kingdom and United States. Street Movement-related channels

at this rank are associated with the *English Defence League* (EDL), a movement opposed to the alleged spread of radical Islamism within the United Kingdom (Goodwin & Ramalingam, 2012). Channels from various international individuals and groups that often describe themselves as “counter-Jihad” can also be observed (Goodwin, 2013). The Conspiracy Theory and non-ER Religion categories appear to be associated with channels based in the United States, where the dominance of these categories at lower rankings (excluding the seed category) suggests that the channels become progressively more U.S.-centric.

The ER Music seed channels usually upload video and audio recordings of high-profile acts associated with the ER. For example, content from bands such as Skrewdriver (United Kingdom) or Landser (Germany) can be found, along with other bands from genres such as Oi!, Rock Against Communism and National Socialist Black Metal (NSBM; Baldauf et al., 2011; Brown, 2004). Given this, the consistent presence of the White Nationalist related category would appear logical. At the same time, we also observe that non-ER Music becomes more evident as k increases, perhaps reflecting the overlap between music genres. For example, someone who is a fan of NSBM is often a fan of other metal music that would not be categorized as ER. For Populist seed channels, the related categories generally appear more diverse. Channels affiliated with political parties can be observed, including the Eurosceptic United Kingdom Independence Party or the British National Party, where the latter is also considered as White Nationalist (Bartlett, Birdwell, & Littler, 2011; Mudde, 2007). Opposition to establishment organizations such as the European Union may be a link to similar opposition within the Patriot and Conspiracy Theory-related categories that are also present (Southern Poverty Law Center, 2013), while also explaining the presence of the non-ER News & Current Affairs category. It should also be mentioned that our definition of Populist is broad and spans multiple themes (Table 1), where a certain amount of disagreement about this category exists (Marlière, 2013).

German Language Categories

A total of 2,766 seed and 177,868 related channel documents were generated from the German language data set, with topics identified using the former. As before, we experimented with values of T_{in} [10,100], and selected $T = 60$, given a similar observation of redundant topic splits for larger values of T . Of these, 33 ER topics (55%), 20 non-ER topics (33.33%), 2 topics that were a combination of ER and non-ER categories (3.33%), and 5 topics based on general terms of a separate language (8.33%) were found. Topic and channel categorization was performed, where we excluded 785 (28.38%) seed and 56,565 (31.8%) related gray sheep channels that could not be categorized, in addition to related channels ranked at $k > 10$ and non-ER seed channels. The remaining 1,123 ER seed and 4,973 related (ER and non-ER, $k \leq 10$) channels were used to calculate the mean related category proportions for each ER category. As previously, the corresponding Freebase topic rankings were determined for the 60 NMF topics (mean percentage of videos having topics was 87% per channel, $\sigma = 17\%$); Table 3 contains a similarly high level of ER topic consistency as shown by the English-language categorization.

Figure 4 contains plots for three ER categories that were selected for in-depth analysis. As seen in Figure 3, the seed category is the dominant related category for all values of k , and the ER-related category presence is consistently stronger than that of the non-ER categories, notwithstanding the increase in diversity. The Populist and Political Party-related categories are prominent for Anti-Islam, given the inclusion of channels affiliated with parties such as the National Democratic Party of Germany, the Pro-Bewegung collective, and the Freedom Party of Austria; all strong opponents of immigration, particularly by Muslims (Baldauf et al., 2011; Bartlett et al., 2011; Goodwin, 2013).

This data set also features many ER Music seed channels that upload recordings of high profile acts, with the main difference being the prominence of the Neo-Nazi related category for all rankings. These recordings and videos, along with other non-music videos uploaded by these channels, often feature recognizable Nazi imagery.

Table 3. Highest ranked German-language ER NMF Topic Categorisation and their Corresponding Freebase Topics, Based on YouTube Annotations of the Sampled Videos.

#Members	Mean Similarity	Category	Top 10 Freebase Topics
14	0.31	Populist	Anonymous, Anti-Counterfeiting Trade Agreement, Occupy movement, Occupy Wall Street, Libya, Scientology, Wiki- Leaks, Muammar al-Gaddafi, Das Ich, Internet censorship
19	0.28	Music, Neo-Nazi, White Nationalist	Frank Rennicke, National Democratic Party of Germany, Landser, Hassgesang, Music, Division Germania, Die Lunikoff Verschwörung, Unseren Toten, Horst Mahler, Projekt Aaskereia
17	0.26	Political Party, Populist, Anti-Islam	Freedom Party of Austria, Heinz-Christian Strache, Vienna, Barbara Rosenkranz, Interview, National Democratic Party of Germany, Stermann & Grisseemann, Natascha Kampusch, Dieter Egger, Ewald Stadler
20	0.24	Music	Bohse Onkelz, Heavy metal, Rock music, Frei.Wild, Viva Los Tioz, Concert, Kategorie C, German rock, Elis, Punk rock
23	0.20	Music, Neo-Nazi, White Nationalist	Sturmwehr, Sleipnir, Nordfront, Stahlgewitter, Kategorie C, Division Germania, Funkenflug, Malchin, Acoustic music, Granite
99	0.20	Political Party, Populist	National Democratic Party of Germany, Udo Pastors, Dresden, Berlin, Anti-Fascist Action, Holger Apfel, Jurgen Rieger, The Left, German People's Union, Frank Rennicke
82	0.18	Anti-Islam	Islam, Muslim, Pierre Vogel, Quran, Documentary, Salafi movement, Muhammad, Thilo Sarrazin, Allah, Jesus Christ
15	0.15	Neo-Nazi	March, Der Hohenfriedberger, Wehrmacht, Prussia, Bad Nenndorf, Prussia's Glory, Military parade, Military, Music, Marching
19	0.14	White Nationalist, Populist, Political Party	British National Party, Nick Griffin, European Union, UK Independence Party, Member of the European Parliament, Nigel Farage, Jobbik, Ashley Mote, European Parliament, Andrew Brons
37	0.14	Anti-Islam, Populist	Islam, Muslim, Iran, Egypt, Israel, Muhammad, Quran, Allah, Geert Wilders, 2011 Egyptian revolution

Note. These NMF topics contain ≥ 10 member seed channels and are ranked using mean pairwise Cosine similarity of the corresponding F-topic vectors. As discussed in Table 1, some of these issues are not exclusive to the extreme right, for example, the Anti-Counterfeiting Trade Agreement (ACTA, considered Populist in an ER context for the purpose of the current work)

Separately, channels that upload videos associated with bands that have alleged ER ties, for example, Bohse Onkelz or Frei.Wild (Baldauf et al., 2011), may explain in part the presence of the non-ER Music category, given the mainstream success of these bands. This may also be explained by material associated with hip-hop acts such as "n'Socialist Soundsystem," which provide an alternative to traditional ER music based on rock and folk. The close relationship with Music is also present for the Neo-Nazi seed category, although further related diversity can be observed. Seed channels featuring footage of German participation in World War II, including speeches by high-ranking members of the Nazi party, are likely to be the source of the White Nationalist and Revisionist-related categories. We can safely assume that the prominence of Music is responsible for the appearance of its non-ER counterpart here.

Aggregate Category View

We conclude our analysis of the ideological bubble created by YouTube’s recommender system by measuring the mean proportions for the seed ER categories as a whole, where the possible aggregated related categories were (1) the same ER category as that of the seed, (2) a different ER category, or (3) a non-ER category. The results for both data sets can be found in Figure 5. As with the individual seed categories, an ER ideological bubble is also clearly identifiable at the aggregate level. Although the increase in diversity for lower k rankings introduces a certain proportion of non-ER categories, this is always outweighed by ER categories, where the seed ER category remains dominant for all values of k . These findings would appear to contrast those of certain prior work where greater related video diversity was observed (Roy et al., 2012). Although we have analyzed related channels rather than individual videos, we might have expected to also find this behavior at both levels. However, it would appear this is not always true, at the very least in the case of ER channels.




Figure 4. German mean category proportions of top k ranked related channels ($k \in [1,3,5,7,10]$) for selected seed ER categories. ER=extreme right; NER = non-extreme right

The data retrieval process described involved following related video links for only one step removed from the corresponding seed video; no further related retrieval was performed for related videos themselves. Of the ER seed channels used in the ideological bubble analysis, we identified those that appeared in top k -related rankings of other ER seed channels, for $k \leq 10$; 6,186 (94%) and 1,056 (94%) such channels, respectively, for the English and German language data sets. These high percentages allude to the presence of cycles within the related channel graph, where retrieving additional data by following related videos for

multiple steps may have been somewhat redundant. They also further emphasize the existence of the ideological bubble.

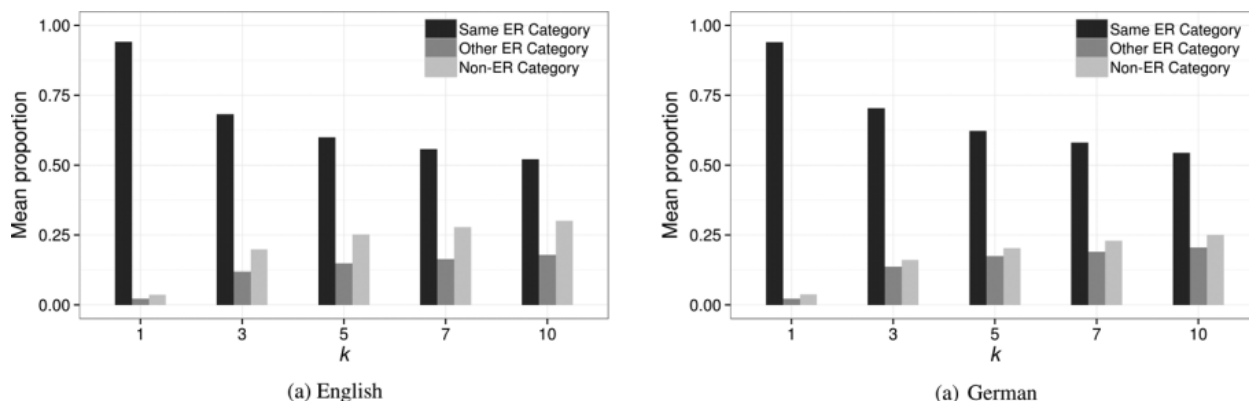


Figure 5. Aggregated mean category proportions of top k ranked related channels ($k \in [1,3,5,7,10]$) for seed ER categories. ER = extreme right; NER = non-extreme right.

Discussion

There is no yet proven relationship between consumption of extremist online content and adoption of extremist ideology (McCants, 2011; Rieger, Frischlich, & Bente, 2013), and some scholars and others remain sceptical of a significant role for the Internet in processes of online radicalization (Benson, 2014; U.K. Home Affairs Committee, 2014, pp. 6–7). There is increasing concern on the part of other scholars, and increasingly also policy makers, that high levels of always-on Internet access and the production and wide dissemination—and hence easy availability—of large amounts of extremist online content may be (violently) radicalizing some of its consumers however (Edwards & Gribbon, 2013). From the producer perspective, this is almost certainly its main purpose. Much recent work in the field of online political extremism has focused upon the relationships among protagonists on specific platforms (Fisher & Prucha, 2014; O’Callaghan et al., 2013a) and the potential outcomes of these actors’ online content production, dissemination, and interaction strategies (Berger & Strathearn, 2013; Carter, Maher, & Neumann, 2014). A majority of this work is thus focused on the convergence of self-communication and mass communication: what Castells (2013, xix) terms “mass self-communication.”

This article takes a different approach focusing on the role of the platform, in this case YouTube, and its sociotechnical infrastructure, specifically its recommender system, and the affordances this mechanism provides for political extremists, in this case, the ER. This is a complex issue, as YouTube’s users are its content providers while the platform itself intervenes via its recommender algorithm to channel specific content to specific users. The immersion of some users in YouTube’s ER spaces is thus a coproduction between the content generated by users and the affordances of YouTube’s recommender architecture (Fiore-Silfvast, 2012, pp. 1967–1968).

The detailed analysis contained herein of how automated social media “recommendation” can result in users being excluded from information that is not aligned with their existing perspective, potentially leading to immersion within an extremist ideological bubble, supports a shift of the almost exclusive focus on users as content creators and protagonists in extremist cyberspaces to also consider platform providers as important actors in these same spaces. This, in turn, suggests that YouTube’s recommender system algorithms are not neutral in their effects but have political articulations:

Together, these algorithms not only help us find information, they provide a means to know what there is to know and how to know it, to participate in social and political discourse, and to familiarize ourselves with the publics in which we participate. They are now a key logic governing the flows of information on which we depend. (Gillespie, 2014, p. 167)

Gillespie goes on to call for close attention to be paid “to where and in what ways the introduction of algorithms into human knowledge practices may have political ramifications.” One of the ways in which he suggests we accomplish this is through reflection upon “the production of calculated publics” or “how the algorithmic presentation of publics back to themselves shapes a public’s sense of itself, and who is best positioned to benefit from that knowledge” (Gillespie, 2014, p. 168). Gillespie supplies the examples of Amazon’s book-buying recommendations invoking a community of readers and Facebook’s “friends of friends” setting transforming a discrete set of users into an “audience”; our focus on YouTube’s creation of an ER milieu is more explicitly political but is nonetheless also an “algorithmically generated group” that “may overlap with, be an inexact approximation of, or have nothing whatsoever to do with the publics that the user sought out” (2014, pp. 188–189), but within which they are then invited to become enmeshed. At the same time, it should be mentioned that YouTube is not the only social media site to be criticized for some of the outcomes of its recommendation practices and the “calculated publics” they “help to constitute and codify . . . [P]ublics that would not otherwise exist except that the algorithm called them into existence” (Gillespie, 2014, p. 189). Twitter’s recommender system has been described by one analyst of violent online jihadism as providing “robust tools . . . to aspiring extremists” and “a running start for users who are interested in pursuing ideologically motivated violence” (Berger, 2013).

It may well be the case that the potential cures for these unintended outcomes are worse than the disease however. Gillespie provides a warning in this respect when he states that “What Twitter claims matters to ‘Americans’ or what ‘Amazon’ says teens read are forms of authoritative knowledge that can and will be invoked by institutions whose aim is to regulate such populations” (2014, p. 190). The regulation of online extremist populations—particularly those espousing violence and/or terrorism—is, for good or ill, now a hot button policy issue. Suggested interventions, besides the legislated takedown of online hate and terrorism content in Germany, the United Kingdom, and elsewhere, have included the setting of tighter standards for acceptable content by platform providers; the insertion of alternative viewpoints into recommended lists associated with certain types of content; the assignment of some videos to the adult category (registered users below a certain age cannot directly view them, and all others users must click their assent to viewing objectionable content); and technical (i.e., algorithm-based) demotion of certain videos containing objectionable content that would otherwise appear on recommended lists (i.e., making otherwise popular content harder to find Berger, 2013; CleanIT, 2013; see also Fiore-Silfvast, 2012). Many of these suggested interventions raise the specter of social media companies policing political thought, which is palatable to neither the companies nor many users, and is especially problematic in the absence of rigorous empirical research that analyzes the Internet’s role in processes of radicalization.

Conclusions and Future Work

YouTube’s position as the most popular video sharing platform has resulted in it playing an important role in the online strategy of the ER. We have proposed a set of categories that may be applied to this YouTube content, based on a review of those found in existing academic studies of the ER’s ideological makeup. Using an NMF-based topic modeling approach, we categorized channels according to this proposed set, permitting the assignment of multiple categories per channel where necessary. The existence of resources such as Freebase allowed us to independently confirm the liability of the categorization where a high level of consistency was observed between our qualitative categorization and YouTube’s automated Freebase annotations. This categorization helped us to identify the existence of an ER ideological bubble, in terms of the extent to which related channels, determined by the videos recommended by YouTube, also belong to ER categories. Despite the increased diversity observed for lower related rankings, this ideological bubble maintains a constant presence. The influence of related rankings on click-through rate (Zhou et al., 2010), coupled with the fact that the YouTube channels in this analysis originated from links posted by ER Twitter

accounts, would suggest that it is possible for a user to be immersed in this content following a short series of clicks.

It might be argued that our findings merely confirm that YouTube's related video recommendation process is working correctly, which is true to a certain extent. Lessig famously stated "code is law" (2006, p. 1); he might equally have said "code is politics." We were concerned with the specificity of YouTube's recommender system in terms of how it works in the world. The article presents a case study of a portion of the politics associated with YouTube's code, along with its potential lived effects; in the absence of any broadly acceptable quick fixes, it seeks to contribute a rigorous evidencing of the already existing political articulations of YouTube's recommender system and to draw attention to the underexplored way in which this may already be influencing political thought and thus potentially also action. The infrastructural affordances of other online platforms differ from those of YouTube; future research into the role of platform providers and their architectures in online extremism could certainly explore whether other popular online platforms thus have similar or different potential effects.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship and/or publication of this article: This research was supported by 2CENTRE, the EU funded Cybercrime Centres of Excellence Network, the EU-funded VOX-Pol Network of Excellence, and Science Foundation Ireland (SFI) under Grant Number SFI/12/RC/2289.

References

- Agosto, D. E., & Hughes-Hassell, S. (2006a). Toward a model of the everyday life information needs of urban teenagers, part 1: Theoretical model. *Journal of the American Society for Information Science and Technology*, 57, 1394–1403.
- Agosto, D. E., & Hughes-Hassell, S. (2006b). Toward a model of the everyday life information needs of urban teenagers, part 2: Empirical model. *Journal of the American Society for Information Science and Technology*, 57, 1418–1426.
- Baldauf, J., Groß, A., Rafael, S., & Wolf, J. (2011) *Zwischen Propaganda und Mimikry: Neonazi-Strategien in Sozialen Netzwerken*. Berlin, Germany: Amadeu Antonio Stiftung.
- Baluja, S., Seth, R., Sivakumar, D., Jing, Y., Yagnik, J., Kumar, S., . . . Aly, M. (2008). Video suggestion and discovery for YouTube: Taking random walks through the view graph. *In Proceedings of 17th international conference on World Wide Web, WWW '08* (pp. 895–904). New York, NY: ACM.
- Bartlett, J., Birdwell, J., & Littler, M. (2011). *The new face of digital populism*. London, England: Demos.
- Bell, M. (2013, March 11). Frei verfügbarer Nazi-Rock: YouTube und die braunen Musikanten. *Der Spiegel*. Retrieved from: <http://www.spiegel.de/netzwelt/netzpolitik/indizierter-nazi-rock-bei-youtube-a-887637.html>
- Benson, D. (2014). Why the internet is not increasing terrorism. *Security Studies*, 23, 293–328.

Berger, J. M. (2013, August 14). Zero Degrees of al Qaeda: How Twitter is supercharging jihadist recruitment. *Foreign Policy*. Retrieved March 12, 2014, from http://www.foreignpolicy.com/articles/2013/08/14/zero_degrees_of_al_qaeda_twitter

Berger, J. M., & Strathearn, B. (2013). Who matters online: Measuring influence, evaluating content and countering violent extremism in online social networks. King's College London: ICSR. Retrieved September 21, 2014, from http://icsr.info/wpcontent/uploads/2013/03/ICSR_Berger-and-Strathearn.pdf

Blee, K. M., & Creasap, K. A. (2010). Conservative and right-wing movements. *Annual Review of Sociology*, 36, 269–286.

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3, 993–1022.

Brown, T. S. (2004). Subcultures, pop music and politics: Skinheads and 'Nazi Rock' in England and Germany. *Journal of Social History*, 38, 157–178.

Bucher, T. (2012). Want to be on the top? Algorithmic power and the threat of invisibility on Facebook. *New Media & Society*, 14, 1164–1180.

Burris, V., Smith, E., & Strahm, A. (2000). White supremacist networks on the internet. *Sociological Focus*, 33, 215–235.

Caiani, M., & Wagemann, C. (2009). Online networks of the Italian and German extreme right. *Information, Communication & Society*, 12, 66–109.

Carter, J., Maher, S., & Neumann, P. (2014). #Greenbirds: Measuring importance and influence in Syrian foreign fighter networks. King's College London: ICSR. Retrieved September 21, 2014, from <http://icsr.info/wp-content/uploads/2014/04/ICSR-Report-Greenbirds-Measuring-Importance-and-Influence-in-Syrian-Foreign-Fighter-Networks.pdf>

Castells, M. (2013). *Communication power*. New York, NY: Oxford University Press.

Chemudugunta, C., Smyth, P., & Steyvers, M. (2006). Modeling general and specific aspects of documents with a probabilistic topic model. In B. Schölkopf, J. Platt, & T. Hofmann (Eds.), *Advances in neural information processing systems 19* (pp. 241–248). Cambridge, MA: Massachusetts Institute of Technology.

Claypool, M., Gokhale, A., Miranda, T., Murnikov, P., Netes, D., & Sartin, M. (1999) Combining content-based and collaborative filters in an online newspaper. In *Proceedings of the ACM SIGIR '99 Workshop on Recommender Systems: Algorithms and Evaluation*. Berkeley, CA.

CleanIT. (2013). Reducing terrorist use of the Internet. Retrieved September 21, 2014, from <http://www.cleanitproject.eu/files/wp-content/uploads/2013/01/Reducing-terrorist-use-of-the-internet.pdf>

Conway, M. (2012). From al-Zarqawi to al-Awlaki: The emergence of the Internet as a new form of violent radical Milieu. *Combating Terrorism Exchange*, 2, 12–22.

Conway, M., & McInerney, L. (2008). Jihadi video and auto-radicalisation: Evidence from an Exploratory YouTube Study. In *Proceedings of the 1st European Conference on Intelligence and Security Informatics* (pp. 108–118). Esbjerg, Denmark: EuroISI '08, Springer-Verlag.

Davidson, J., Liebald, B., Liu, J., Nandy, P., Van Vleet, T., Gargi, U., . . . Sampath, D. (2010). The YouTube video recommendation system. In *Proceedings of the 4th ACM conference on Recommender Systems* (pp. 293–296).

RecSys '10. New York, NY: ACM.

Edwards, C., & Gribbon, L. (2013). Pathways to violent extremism in the digital era. *RUSI Journal*, 158, 40–47.

Europol. (2011). *TE-SAT 2011: EU terrorism situation and trend report*. The Hague, The Netherlands: European Police Office. Retrieved March 12, 2014, from <https://www.europol.europa.eu/sites/default/files/publications/te-sat2011.pdf>

Figueiredo, F., Benevenuto, F., & Almeida, J. M. (2011). The Tube over time: Characterizing popularity growth of YouTube videos. In *Proceedings of the 4th ACM international conference on Web search and data mining, WSDM '11* (pp. 745–754). New York, NY.

Filippova, K., & Hall, K. B. (2011). Improved video categorization from text metadata and user comments. *Proceedings of the 34th International ACM SIGIR Conference on Research and development in Information Retrieval* (pp. 835–842). New York, NY: ACM, SIGIR '11.

Fiore-Silfvast, B. (2012). User-generated warfare: A case of converging wartime information networks and coproductive regulation on YouTube. *International Journal of Communication*, 6, 1965–1988. Retrieved from September 23, 2014, <http://ijoc.org/index.php/ijoc/article/viewFile/1436/774>

Fisher, A., & Prucha, N. (2014). The call-up: The roots of a resilient and persistent Jihadist presence on Twitter. *Combating Terrorism Exchange*, 4, 73–88.

Gerstenfeld, P. B., Grant, D. R., & Chiang, C. P. (2003). Hate online: A content analysis of extremist Internet sites. *Analyses of Social Issues and Public Policy*, 3, 29–44.

Gillespie, T. (2003). The stories that tools tell. In J. Caldwell & A. Everett (Eds.), *New media: Theses on convergence media and digital reproduction*. London, England: Routledge.

Gillespie, T. (2010). The politics of 'platforms.' *New Media & Society*, 12, 347–364.

Gillespie, T. (2014). The relevance of algorithms. In T. Gillespie, P. Boczkowski, & K. Foot (Eds.), *Media technologies: Essays on communication, materiality, and society*. Cambridge, MA: MIT Press.

Goodwin, M. (2013). *The roots of extremism: The English defence league and the counter-Jihad challenge*. London, England: Chatham House.

Goodwin, M., & Ramalingam, V. (2012). *The new radical right: Violent and non-violent movements in Europe*. London, England: Institute for Strategic Dialogue.

Greene, D., & Cunningham, P. (2013). Producing a unified graph representation from multiple social network views. In *Proceedings of the 5th Annual ACM Web Science Conference* (pp. 118–121). WebSci '13, ACM, New York, NY.

Hannon, J., Bennett, M., & Smyth, B. (2010). Recommending Twitter users to follow using content and collaborative filtering approaches. In *Proceedings of the 4th ACM Conference on Recommender Systems (RecSys'10)* (pp. 199–206). New York, NY: ACM Press.

Ito, M., Baumer, S., Bittanti, M., boyd, d., Cody, R., Herr-Stephenson, B., ... Tripp, L. (2010). *Hanging out, messing around, and geeking out: Kids living and learning with new media*. Cambridge, MA.: MIT Press.

Langlois, G., & Elmer, G. (2013). The research politics of social media platforms. *Culture Machine*, 14, 1–17.

Lee, D. D., & Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature*,

401, 788–791.

Lessig, L. (2006). *Code: Version 2.0*. New York, NY: Basic Books.

Livingstone, S., & Haddon, L. (2009). EU Kids Online: Final Report. Retrieved [http://www.lse.ac.uk/media@lse/research/EUKidsOnline/EU%20Kids%20I%20\(2006-9\)/EU%20Kids%20Online%20I%20Reports/EUKidsOnlineFinalReport.pdf](http://www.lse.ac.uk/media@lse/research/EUKidsOnline/EU%20Kids%20I%20(2006-9)/EU%20Kids%20Online%20I%20Reports/EUKidsOnlineFinalReport.pdf)

Livingstone, S., Ólafsson, K., O'Neill, B., & Donoso, V. (2012). Towards a better internet for children. Retrieved March 12, 2014 from <http://www.lse.ac.uk/media@lse/research/EUKidsOnline/EU%20Kids%20III/Reports/EUKidsOnlineReportfortheCEOCoalition.pdf>

McCants, W. (2011). Testimony, U.S. house of representatives, subcommittee on counterterrorism and intelligence, Jihadist use of social media: How to prevent terrorism and preserve innovation, 6 December. Retrieved September 21, 2014, from <http://homeland.house.gov/sites/homeland.house.gov/files/Testimony%20McCants.pdf>

Marlière, P. (2013). Populism and the enchanted world of 'moderate politics'. Retrieved March 12, 2014, from <https://www.opendemocracy.net/can-europe-make-it/philippe-marliere/populism-and-enchanted-world-of-moderate-politics>

Mudde, C. (2007). *Populist radical right parties in Europe*. Cambridge, England: Cambridge University Press.

O'Callaghan, D., Greene, D., Conway, M., Carthy, J., & Cunningham, P. (2013a). An analysis of interactions within and between extreme right communities in social media. *Ubiquitous Social Media Analysis, 8329*, 88–107.

O'Callaghan, D., Greene, D., Conway, M., Carthy, J., & Cunningham, P. (2013b). Uncovering the wider structure of extreme right communities spanning popular online networks. In *Proceedings of the 5th ACM Web Science Conference* (pp. 276–285). New York, NY: WebSci '13.

Pariser, E. (2011). *The filter bubble: What the Internet is hiding from you*. London, England: Penguin.

Rieger, D., Frischlich, L., & Bente, G. (2013). *Propaganda 2.0: Psychological effects of right-wing and Islamic extremist Internet videos*. Cologne, Germany: Wolters Kluwer.

Roy, S. D., Mei, T., Zeng, W., & Li, S. (2012). SocialTransfer: cross-domain transfer learning from social streams for media applications. In *Proceedings of the 20th ACM international conference on Multimedia* (pp. 649–658). MM '12, ACM. New York, NY.

Simonet, V. (2013). Classifying YouTube channels: A practical system. In *Proceedings of the 2nd International Workshop on Web of Linked Entities (WoLE 2013)*, (pp. 1295–1304). ACM.

Slack, J. D., & Wise, J. M. (2007). *Culture and technology: A primer*. New York, NY: Peter Lang Publishing.
Southern Poverty Law Center. (2013). Ideology. Retrieved March 12, 2014, <http://www.splcenter.org/get-informed/intelligence-files/ideology>

Steyvers, M., & Griffiths, T. (2006). Probabilistic topic models. In T. Landauer, D. Mcnamara, S. Dennis, & W. Kintsch (Eds.), *Latent semantic analysis: A road to meaning*. Mahwah, NJ: Laurence Erlbaum.

Sunstein, C. R. (2001). *Republic.com*. Princeton, NJ: Princeton University Press.

Tateo, L. (2005). The Italian extreme right on-line network: An exploratory study using an integrated social

network analysis and content analysis approach. *Journal of Computer-Mediated Communication*, 10. Retrieved from: <http://onlinelibrary.wiley.com/doi/10.1111/j.1083-6101.2005.tb00247.x/abstract> - here it states the page number was 00.

U.K. Home Affairs Committee. (2014). *Counter-terrorism: Seventeenth report of session 2013–14*. London, England: The Stationery Office. Retrieved September 21, 2014, from <http://www.publications.parliament.uk/pa/cm201314/cmselect/cmhaff/231/231.pdf>

U.K. Home Office. (2011) *Contest: The United Kingdom's strategy for countering terrorism*. Norwich, England: The Stationery Office. Retrieved March 12, 2014, https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/97995/strategy-contest.pdf

YouTube Team. (2008). Dialogue with Sen. Lieberman on terrorism videos. Broadcasting Ourselves: *The YouTube Blog* 19 May. Retrieved March 12, 2014, from <http://googlepublicpolicy.blogspot.ie/2008/05/dialogue-with-sen-lieberman-on.html>

Zhou, R., Khemmarat, S., & Gao, L. (2010). The impact of YouTube recommendation system on video views. *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement* (pp.404–410), ACM, New York, NY.

Author Biographies

Derek O'Callaghan is a PhD student in the School of Computer Science & Informatics, University College Dublin. His research focuses on social media analytics, community finding, and text mining, in the analysis of online extremist activity; email: derek.ocallaghan@ucd.ie

Derek Greene is a lecturer in the School of Computer Science & Informatics, University College Dublin, and a researcher at the Insight Centre for Data Analytics. His current research focuses on scalable machine learning, network analysis, and recommender systems; email: derek.greene@ucd.ie

Maura Conway is a senior lecturer in the School of Law & Government, Dublin City University. Her research interests are in the intersections of violent political extremism and the Internet. She is the principal investigator on the European Union FP-7-funded VOX-Pol project on violent online political extremism; email: maura.conway@dcu.ie

Joe Carthy is the college principal and Dean of Science at University College Dublin. He was previously Head of the UCD School of Computer Science & Informatics. He is active in two research areas: cybercrime investigation and forensic computing and information retrieval; email: joe.carthy@ucd.ie

Pádraig Cunningham is a professor of knowledge and data engineering in the School of Computer Science & Informatics, University College Dublin. His current research focus is on the analysis of graph and network data and on the use of machine learning techniques in processing high-dimension data; email: padraig.cunningham@ucd.ie