

Commenting on YouTube Videos: From Guatemalan Rock to El Big Bang¹

Mike Thelwall, Pardeep Sud, Statistical Cybermetrics Research Group, School of Technology, University of Wolverhampton, Wulfruna Street, Wolverhampton WV1 1SB, UK.

Farida Vis, Department of Media and Communication, Leicester University, University Road Leicester LE1 7RH, UK.

YouTube is one of the world's most popular web sites and hosts numerous amateur and professional videos. Comments on these videos may be researched to give insights into audience reactions to important issues or particular videos. Yet little is known about YouTube discussions in general: how frequent they are, who typically participates and the role of sentiment. This article fills this gap through an analysis of large samples of text comments on YouTube videos. The results identify patterns and give some benchmarks against which future YouTube research into individual videos can be compared. For instance, the typical YouTube comment was mildly positive, was posted by a 29 year old male, and contained 58 characters. About 23.4% of comments in the complete comment sets were replies to previous comments. There was no typical density of discussion on YouTube videos in the sense of the proportion of replies to other comments: videos with few replies and with many replies were both common. The YouTube audience engaged with each other disproportionately when making negative comments, however; positive comments elicited few replies. The biggest trigger of discussion seemed to be religion, whereas the videos attracting the least discussion were predominantly from the Music, Comedy and How to & Style categories. This suggests different audience uses for YouTube: from passive entertainment to active debating.

Introduction

The online video sharing web site YouTube, which was originally created in February 2005 to help people share videos of well-known events (Hopkins, 2006), has rapidly grown to be a cultural phenomenon for its mass user-base. It seems to have attracted little social science research compared to general social network sites (SNSs) despite apparently being the third most popular web site globally according to Alexa (<http://www.alexa.com/topsites>, as of June 3, 2011). YouTube is also interesting as a site driven to a large extent by freely-contributed content, with uploaders being motivated and rewarded by viewers' attention rather than money (Huberman, Romero, & Wu, 2009). In June 2009, 69% of US internet users had accessed videos and 14% had posted videos (females as much as males), although not necessarily on YouTube (Purcell, 2010). The relative lack of social science research may be because a common activity is watching TV-like content, such as music videos and TV shows (Waldfoegel, 2009). Nevertheless, YouTube makes it easy for people with a video recording device and internet connection to publish their own videos and some of these amateur videos have attracted tens of millions of hits (e.g., *Charlie bit my finger - again*², with 283,629,150 views by February 22, 2011, and *Chinese Backstreet Boys - That Way*³, with 13,052,790 views by February 22, 2011) or a moderate number of hits, but still a large audience for an amateur production (e.g., *Lynne and Tessa*⁴, with 52,081 views by February 22, 2011). Moreover, the convenience of YouTube seems to be widely used for semi-professional video productions, from organisations' *About us* or *Welcome* videos to recordings of lectures or demonstrations of how to do something (e.g., *Natural Looking Makeup Tutorial*⁵, with 5,225,414 views by February 22, 2011) and professional or amateur videos about illnesses (Lo, Esser, & Gordon, 2010).

YouTube and other online video services have become part of the political process in some countries, such as the US (Gueorguieva, 2008) and South Korea (Hang & Yun, 2008; Im, 2010),

¹ This is a preprint of an article to be published in the Journal of the American Society for Information Science and Technology © copyright 2011 John Wiley & Sons, Inc.

² http://www.youtube.com/watch?v=_OB1gSz8sSM

³ <http://www.youtube.com/watch?v=N2rZxCrb7iU>

⁴ http://www.youtube.com/watch?v=mN2_lzWGaCg

⁵ <http://www.youtube.com/watch?v=OB8nfJCOIeE>

although their influence may be typically small (Baumgartner & Morris, 2010). Occasionally, however, YouTube videos can have a significant impact on the outside world. One music video by a dissatisfied customer apparently cost an airline 10% of its share price (Ayres, 2009) and a video of the death of Iranian protester Neda initially spread on YouTube and Facebook (Van Langendonck, 2009) and triggered international media coverage. There is also some evidence that prominent news events are reflected by increased associated YouTube video posting (Sykora, & Panek, 2009a) and even that stock market movements may have associated YouTube posting trends (Sykora, & Panek, 2009b). One interesting feature of YouTube is its interactivity because viewers can post video responses or text comments after watching a video. Despite the research potential of such public audience reactions (e.g., Losh, 2008) and the possible value of the feedback to the video owners (e.g., Fauconnier, 2011), there is no systematic research into how they work in the sense of how common they are, who takes part and which issues trigger the most and the least debate.

Most YouTube research seems to take a humanities perspective, typically investigating one video genre and focusing on the purpose and/or reception of that genre (e.g., childbirth, coming out) or particular topics (Thorson, Ekdale, Borah, Namkoong, & Shah, 2010), types of information (Steinberg et al., 2010) or potential threats to society from the information disseminated (Lewis, Heath, St Denis, & Noble, 2011). This has shown that amateur YouTube videos fulfil a wide variety of social needs and may evoke a more personal relationship between the viewer and viewed compared to other online publishing. There have also been some large-scale quantitative analyses of YouTube (reviewed below) but none have focused on audience reactions in the form of comment-based discussions.

This article addresses one aspect of YouTube videos: the textual comments posted in response to them. When someone views a video, they can respond or interact in four ways unless the owner has disabled the features: by rating the video or a comment as good or bad, by posting a video response or by posting a comment about the video to the video page. A US survey from early 2007 found that 13% of users watching online videos had posted comments about them (Madden, 2007) and the data collected in the current paper suggests that there is one comment for every 204 views of a YouTube video that attracts at least one comment – 0.5% of viewers leave a comment. This article focuses on the section of the YouTube audience that writes comments and the extent to which these comments become debates. The goal is to generate baseline statistics so that future researchers can tell whether the videos that they are investigating are typical or unusual. Although comments are a relatively minor aspect of YouTube, they are socially significant because of YouTube's mass user base. Whilst the main quantitative evidence about video popularity comes from total viewer numbers, the number of positive and negative ratings and the number of times that a video has been favoured, the focus on commenters may give deeper insights into the YouTube audience and the second focus on debates may give insights into what is controversial or triggers discussion in other ways. Video comments are ignored because they would require a different kind of analysis and would presumably be created by a different kind of viewer. Nevertheless, since a small proportion of viewers comment on a video, the extent to which comments can give audience insights is limited. Although anybody can watch YouTube videos, they must register with the site in order to post a comment. As part of this registration process they may volunteer personal information such as age, gender and location (and may lie, of course) and this information is accessible to researchers either on the YouTube web site or via the YouTube API (<http://code.google.com/apis/youtube/overview.html>, accessed February 22, 2011). No information is available about viewers that do not comment, although YouTube gives broad viewer statistics on some videos via a "Show video statistics" button.

Background

This section introduces the theoretical and factual background in terms of research into online discussions and into uses of YouTube.

Online discussions

Many studies have investigated the extent to which online communication differs from offline communication and differs between online contexts (Herring, 2002). In contrast to typical face-to-face communication, online communication may be anonymous, textual, asynchronous, remote, permanent

and/or very public, although some online forms can be none of these. This review focuses on contexts that have at least one of the above properties, since public comments in YouTube have them all.

YouTube commenters can choose to be anonymous because even though they must register an identity to comment, they may use a pseudonym and this seems to be the norm (from a visual inspection of the data gathered for the current research). Anonymity seems to partly free participants from social norms, perhaps because of the practical impossibility of imposing social or other sanctions on anonymous users in most contexts (Friedman, Khan, & Howe, 2000). This may lead to antisocial behaviour, such as flaming (Alonzo & Aiken, 2004), but other factors may provide an alternative explanation; see below. In practice, YouTube commenters may choose a pseudonym that their friends would be aware of, such as their nicknames. This would be likely to make their offline identities transparent to their friends but hidden from strangers.

YouTube comments are textual and much research has investigated the limitations and peculiarities of electronic text. Early studies were particularly concerned that the absence of the non-verbal channel in textual communication would lead to widespread misunderstandings, particularly in short message formats, such as mobile phone texting (Walther & Parks, 2002). In response, however, a number of conventions have emerged to express sentiment in short informal text, such as emoticons and deliberate non-standard spellings (e.g., Derks, Bos, & von Grumbkow, 2008). In open forums, various conventions have also arisen to signify to whom a message is directed, such as the @ symbol, and its topic (via an embedded hashtag or a meta-tag), and there is evidence that the @ symbol is extensively used for discussions in Twitter (e.g., Java, Song, Finin, & Tseng, 2007; Kwak, Lee, Park, & Moon, 2010), where hashtags and the @ convention probably emerged.

Asynchronous online discussions, such as those via YouTube comments, are those where there may be delays between contributors, perhaps because they live in different time zones or log on at different times of day. Asynchronous communication seems likely to defuse emotions in online discussions since emotions are, by their nature, short term events (although moods last longer) (Cornelius, 1996).

An important issue for this paper is the types of topics that are discussed online most and the triggers of discussions. An analysis of dialogs in the social network site MySpace found most exchanges to be friendly and sociable, often performing the function of keeping in touch with friends and acquaintances (Thelwall & Wilkinson, 2010). A number of projects have shed light on the dynamics of online discussions in terms of what triggers and sustains contributions, what kind of people contribute at different stages, and what the typical structures of discussions are. One study, of a news forum, has found that negativity sustains discussions because the longest threads tended to have negative sentiments expressed at their beginning (Chmiel et al., 2011). A similar result has been found in a case study in Twitter (Naveed, Gottron, Kunegis, & Alhadi, 2011). Possibly related to this, longer discussions in a Polish forum were found to be associated with controversial topics (Sobkowicz & Sobkowicz, 2010).

YouTube audiences and discussion topics

Although there have been some large-scale quantitative investigations into YouTube (Ding et al., 2009; Gill, Arlitt, Li, & Mahanti, 2007), few have focused on discussions in comments. Most YouTube research seems to be small-scale and qualitative, able to give insights into how discussions can occur around videos without giving broad overall patterns of use. An exception is the discovery that there are patterns in user types that can be used to predict users' likely behaviours (Maia, Almeida, & Almeida, 2008).

For online video watching in general, a study of US internet users in 2009 found that 50% of adults had watched funny videos, 38% had watched educational videos, 32% had watched TV shows or movies and 20% had viewed political videos (Purcell, 2010). Nevertheless, it seems likely that people may watch a particular category much more often than another, so these percentages may not be representative of what is typically watched online.

In terms of common content categories in YouTube, music videos are a significant presence in YouTube, probably accounting for about a quarter of videos, at least in April 2007, with entertainment, comedy and sports categories all accounting for very approximately 10% of posted

videos each (Cheng, Liu, & Dale, in press). Perhaps related to this, most videos are quite short, with the modal length being 20-40 seconds and the majority being under 4 minutes (Cheng et al., in press).

From the popular categories, the sports genre is perhaps the most obvious source of controversial content. Sports videos often show highlights of competitions as well as controversial and unusual occurrences (Stauff, 2009). Moreover, a competition has winners and losers, with supporters of both sides, and so it seems reasonable to expect arguments between opposing sides and perhaps performance dissections from supporters – with these dissections drawing upon a rich culture of history and information use in media-led sports discussions (Stauff, 2009).

In contrast to highly mediated content, another study found that amateur videos are capable of attracting a real audience, albeit a small one. For example, 60% of videos are watched at least 10 times during the first day in which they are posted (Cha, Kwak, Rodriguez, Ahn, & Moon, 2009). Nevertheless, a previous study suggested that 10% of videos account for about 80% of views (Cha, Kwak, Rodriguez, Ahn, & Moon, 2007). The first study also showed that videos that did not attract many viewers within the first few days of publication were unlikely to grow an audience later on (Cha et al., 2009). Some small-scale studies have asserted that amateur YouTube videos have a personal and intimate nature, often being filmed in a bedroom or at home (Molyneaux, O'Donnell, Gibson, & Singer, 2008). This may make it easy for viewers to empathise with authors, and hence it would be reasonable to expect predominantly positive comments (e.g., Lazzara, 2010). For example, the “coming out” video seems to be a recognised genre, with many preferring to come out online before offline, presumably in the expectation of a better response, perhaps from a targeted set of friends informed about the video location, or at least increased personal safety (Alexander & Losh, 2010).

Like social network sites, such as Facebook, YouTube has a Friend network and in January 2007 just under 80% of Friend-like subscriber connections were reciprocal but users had only an average of 4 connections each and were members of an average of 0.25 groups (Mislove, Marcon, Gummadi, Druschel, & Bhattacharjee, 2007). Whilst the Friend network may be irrelevant for many or most discussions, it seems likely to be relevant for discussions of personal videos because many of these would only be interesting to people knowing those filmed (Lange, 2009). The Friend network can also be relevant for other topics, however, such as politics. For instance, an investigation into video and textual responses to the controversial anti-Islam Fitna video found that a core of discussion contributors (i.e. commenters) were connected to each other as YouTube Friends or had shared interests, as evidenced by common YouTube channel subscriptions (van Zoonen, Mihelj, & Vis, in press). This shows that comment contributions may draw upon a network of known individuals, even when the commented video is of widespread interest (e.g., in the news). A study of the YouTube network, based upon a crawl of Friend connections, found that people tended to connect to others producing similar content, as measured by tags added to videos by their authors (Paolillo, 2008).

Factors impacting behaviour in YouTube discussions

YouTube has the technical capacity to host debates via comments or video replies. Nevertheless, YouTube “is not primarily designed for collaborative or collective participation”, although it occurs for a minority of users (Burgess & Green, 2009, p. 63, see also Chapter 4). One way in which YouTube can trigger collective action is by viewers creating videos in response to others. In comparison to commenting, this process seems to be too slow to generate significant debates, however. A study of frequently imitated videos found them to have “A focus on ordinary people, flawed masculinity, humor, simplicity, repetitiveness, and whimsical content” (Shifman, in press). Some studies demonstrate that YouTube hosts significant commentary, if not debate, for some important issues, however. One example is the Fitna film of Dutch politician Geert Wilders, mentioned above, which triggered video responses and extensive commenting in YouTube (van Zoonen et al., in press). The extent of the reaction prompted the claim that YouTube had become a mainstream venue for publishing opinions about this issue (van Zoonen, Vis, & Mihelj, 2010). The Fitna case may be somewhat unusual, however, since the film was initially released as an online video (although not on YouTube) and, therefore has a natural fit with YouTube. In contrast, typical news stories might be more suited to debates in political blogs or discussion forums or via news web sites.

There has been interest in the potential for the internet to facilitate exchanges of views amongst citizens: a type of “public sphere” (Habermas, 1991) for political debates (Castells, 2008).

The blogosphere seems to be the most logical place for serious discussions because blog posts can be as long as the author chooses (unlike Twitter and YouTube comments) and can connect to other posts (e.g., Tremayne, Zheng, Lee, & Jeong, 2006). In contrast, some have argued that the diversity of content on the internet allows people to choose to only view material that they agree with, hence avoiding any genuine debate or alternative perspectives (Sunstein, 2007). Perhaps in alignment with the latter point, a study of YouTube videos of Atlantic Canada found little evidence of viewers engaging in discussions online, although most viewers talked offline about the videos that they had seen (Milliken, Gibson, O'Donnell, & Singer, 2008; Milliken, Gibson, & O'Donnell, 2008). This is relevant to the uses and gratifications theory (Blumler & Katz, 1974), which claims that people do not always consume media passively but often use it for their own goals – such as for future conversation topics. Overall, however, this shows that the impact of internet videos may be wider than apparent from the comments on them.

The current study is concerned with public videos in YouTube and the comments on these, if any, will also be public. Note that YouTube comments are text only (e.g., no HTML, URLs or embedded images). In principle, anyone with web access can view any public YouTube comment and in practice commenters can expect their messages to be read by at least some unknown people. This may make users more cautious about what they write, particularly if their YouTube accounts are not anonymous. Nevertheless, YouTube users often seem to treat their privacy casually (Lange, 2007b) and so the public nature of comments may not greatly restrict expressiveness. Another factor that may induce caution is the relative permanence of YouTube comments. Although they will disappear if the hosting video is deleted, this may not happen and the comments could become permanently available on the web. Nevertheless, comments on unpopular videos are likely to be rarely read and comments on popular videos are also likely to become rarely read as they are replaced by newer comments at the top of the list.

Participants in YouTube discussions may be geographically remote. This remoteness means that participants may be more mixed in terms of culture than is common offline, which may lead to misunderstandings. Participants may also mix outside of their normal social circle, in terms of age and gender, which may cause further misunderstandings. Related to this, YouTube use can be regarded very differently by participants. Some may regard themselves as members of the network and behave accordingly, such as following politeness rules of behaviour, whereas others may regard themselves as visitors or regard YouTube as an anarchic environment (Lange, 2008).

Some information is available on YouTube comments and commenters. One important factor concerning antagonisms between commenters is that YouTube users have differing beliefs about acceptable behaviour, which causes friction when a person writes something that they consider acceptable but that antagonises others. The paper also argues that this is more likely to be the primary cause of antagonisms in YouTube than anonymity (Lange, 2007a). A large-scale study using 756 popular queries to generate 67,290 videos with 6.1 million comments has investigated the role of sentiment in categories and the ratings of comments (i.e., the extent to which YouTube users rate a comment as good or bad), finding that ratings were predominantly positive. This study also categorised comments with probabilities to be positive, negative or neutral using a simple machine learning approach based upon a sentiment word list and found that negative comments tended to be disliked and positive comments tended to be liked (Siersdorfer, Chelaru, Nejd, & Pedro, 2010). Moreover, the average sentiment of comments and their average ratings varied by video category, with the Music category having the highest ratings and most positive comments. The three categories with the most negative comments were Shows, Nonprofits – Activism, and Comedy. In two of these cases the content could have been often thought unfunny or not entertaining but in the political example, it could be that people disagreed with the content of the video instead (Siersdorfer et al., 2010). Another study investigated only epilepsy-related videos but found that official videos were less likely to attract comments and empathy than amateur videos (Lo et al., 2010). This seems likely to be true for other types of video too because the audience may feel closer to amateur producers.

Research questions

The goal of this study is to generate descriptive statistics about YouTube comments, and particularly about discussions via YouTube comments. Although there have been some quantitative and

qualitative studies of YouTube, not enough is known about its uses in general to be able to formulate hypotheses about why discussions might occur. For example, the following all seem to be reasonable causes of discussions but there is insufficient evidence to make a credible claim that one is likely to be dominant or that other causes are less likely: discussions are triggered by disagreements about controversial topics; discussions occur to identify unknown facts (e.g., who appeared in a video); discussions are purely social (phatic); discussions are mainly offers of social support. Hence, no prior hypotheses are made about the main *causes* of discussions. Instead, the following general exploratory research questions drive the study.

- What are the typical characteristics of authors of comments on YouTube videos?
- What are the typical characteristics of comments on YouTube videos?
- What are the key topics and factors that trigger discussions on YouTube videos?

Data and methods

A large sample of YouTube video comments and commenters was needed to find typical characteristics. Although it is possible in theory to randomly sample YouTube videos because video IDs are assigned at random (Cheng et al., in press) there is no exhaustive list or a searchable ID space, which makes random sampling difficult. We therefore adapted a method to generate a large sample of videos from which a small test set could be randomly selected (Siersdorfer et al., 2010). For this, we extracted a list of 65,536 terms from a set of predominantly English blogs and RSS feeds used for other purposes. The variety in this source should ensure that unpopular videos are retrieved in addition to popular ones. We used Webometric Analyst (<http://lexiurl.wlv.ac.uk>) to submit these terms individually as single word queries to YouTube via its applications programming interface (API). Webometric Analyst selected one video at random for each search and downloaded its comments, again using the YouTube API. Each query returned a list of up to 1,000 matching video IDs, with 40,997 queries returning at least one comment. We then retrieved the first up to 1,000 comments from each video in the list of 40,997, again from Webometric Analyst using the YouTube API, and identified whether each comment was a reply to a previous comment in the same set, as flagged in the data returned from the API. This information together formed our *comments sample*. Note that this is not a random sample of YouTube due to the English bias in the origins of the word list. Others have used alternative strategies to gather YouTube samples, such as crawling the site using Friend connections (Mislove et al., 2007). This method produces lists of users rather than lists of videos, however, and is very resource-intensive because it needs to cover a high proportion of the network of users to avoid biases caused by the snowball-type method used. The previous similar method that used Google's Zeitgeist for the query terms is also undesirable for the current paper as it focuses on popular topics. The method used here is a compromise and somewhat hybrid because it produces unknown proportions of popular and unpopular videos and so matches neither the videos viewed by users nor the videos posted by users. Nevertheless, it seems to be a reasonable choice for the task.

There were 1,605 videos in the comments sample with 999 or 1,000 comments returned by the API. These probably all had over 1,000 comments, but the number returned was truncated to about 1,000 due to the API limit of 1,000 comments returned per video. For instance, one of the videos with 1000 comments returned had an estimated 366,878 comments in total, with the most recent 1,000 returned by the API. In order to study complete discussions, a second data set was extracted from the comments sample by removing all videos with 999 or 1,000 comments returned by the API - i.e., the videos with incomplete comment sets. This resulted in 39,392 videos. One comment was selected at random from each video and information about the commenter extracted from the YouTube API using Webometric Analyst. The resulting information for 38,628 commenters formed our *commenters sample*. In order to study the extent to which debates occurred in the comments, videos with only 1 comment were also removed as these could not be a discussion. The remaining 35,347 videos formed our *complete comment sets* sample. Note that the exclusion of the 4% of videos with 999 or 1000 comments is a limitation of the research. The overall results should not be greatly impacted by the 4% removal because the percentage removed is so small, however, with the exception of the mean comments per video, which is reported below for reference. Accurate statistics about reply density cannot be calculated from these because the data is incomplete and because comments in the first

1000 may be replies to comments outside the first 1000. To give an extreme but plausible example, many of the first 1000 comments on a popular video may be rejoinders to a particularly offensive recent comment, with few earlier comments being replies. The discussion density of the most recent 1000 comments would therefore be much higher than for the entire discussion.

The samples were processed to extract summary information for the key data returned by the YouTube API. This is a data-driven or information-centred (Thelwall, Wouters, & Fry, 2008) approach since it exploits the data available from YouTube rather than starting with a theoretically-driven set of requirements for information about YouTube and devising methods to obtain the information. The methods for each summary, when not obvious, are described in the results section.

Sentiment strengths for comments were measured using SentiStrength (Thelwall, Buckley, Paltoglou, Cai, & Kappas, 2010, downloaded from <http://sentistrength.wlv.ac.uk>), which is sentiment analysis software that is designed to measure sentiment strengths in short informal English text - predominantly the type in the YouTube comment sample. SentiStrength works mainly by identifying sentiment-related words in a text (e.g., hate) and using all the sentiment words found in a scoring function to predict the overall sentiment of the text. Its accuracy was assessed on a set of 3407 human-coded YouTube comments and it gave a Spearman correlation of 0.583 for positive sentiment and 0.518 for negative sentiment, indicating that it approximates human levels of accuracy at detecting sentiment strength (Thelwall et al., 2010). To filter out non-English comments, each YouTube video was discarded for the sentiment analysis unless at least one comment contained at least one common and fairly distinctive English word (e.g., the) and no comments contained any word from a small set of distinctive non-English words (e.g., el, la, le, al, das, ja). This resulted in 1,242,885 comments on 9,592 videos. These were copied into a single text file (one comment per line) and fed to SentiStrength for sentiment strength classification.

For the third research question, a logical and easily identified proxy for the extent to which comments form a discussion is to calculate the proportion of comments that are recorded as replies to other comments. Whilst it is possible to discuss in YouTube without using the formal reply function when posting a new comment, it is difficult to automatically identify such informal replies because of the need for complex natural language processing techniques to identify inter-comment linguistic references. Hence, comments were assumed to be participating in a discussion only if they were replies to previous comments. This information should therefore be treated as a lower bound for the amount of discussion. Using terminology from social network analysis (SNA), each discussion can be viewed as a network with the nodes being the comments and two nodes being connected if one comment is a reply to another. The density of this network, irrespective of its size, therefore represents the intensity of the discussion: it is the number of connected pairs of nodes divided by the total number of possibly connected pairs of nodes [$\#replies/(\#comments-1)$]. Note that this is the spirit but not the formula for the standard SNA density metric (Wasserman & Faust, 1994) – the standard formula is inappropriate because comments can reply to a maximum of one other comment.

Results and discussion

The results reported below are organised separately for each of the three data samples. The first three subsections primarily include basic findings whereas the final subsection includes a more detailed analysis.

Individual commenters

This subsection gives broad summary statistics about commenters to serve as context for this study and for future investigations into YouTube commenting.

Age and gender The commenters sample was analysed for reported age and gender. Of these commenters, 37,533 (97.2%) recorded a gender, with almost three quarters (72.2%) being male. Figure 1 displays the overall distribution of commenter ages for the 33,923 (87.8%) that declared an age. The most common age was 20, the median was 25 and the mean was 29.3 years old. Almost 1% of commenters reported an age of 109, suggesting misrepresentation, and this may also be the explanation for the outlying bars at round numbers: ages 30 and 20 and, to a lesser extent, 40. Nevertheless, YouTube commenters seem to be young on average but, even allowing for age falsification, are probably not predominantly teenagers. Males were an average of 2.3 years older

than females (mean 29.9 compared to 27.6 for females; Mann-Whitney U test for rank differences, $p=0.000$). Older members also tended to write longer comments (Spearman's $\rho = 0.144$, $p=0.000$) but comment length was unrelated to gender (Mann-Whitney U test, $p=0.056$).




Figure 1. Self-reported commenter ages for a random comment from each video retrieved (one selected at random per search) with at least one comment.

Location Most commenters declared a country location, and were predominantly from the USA, as Table 1 shows, but almost two thirds were from elsewhere in the world. This is broadly in line with YouTube press information from June 2011 reporting 70% of “traffic” to originate from outside the US (http://www.youtube.com/t/press_statistics, accessed June 17, 2011), given the English bias of the data used here. Despite the English bias of the original list, most countries in the table do not have English as a dominant language, even though some majority English-speaking nations, like Ireland (0.6%) and New Zealand (0.4%), are not included in the list. In total, 51.3% of comments derive from nations where English is the dominant language so probably about half of the commenters, overall, are native English speakers, allowing for a minority of non-native English speakers in these countries. Note that the YouTube audience is partly constrained by attempts to block it from various countries, most notably China (Sommerville, 2009).

Table 1. Declared location of 37,595 commenters - countries with at least 1.0% of the commenters are shown.

<i>Country</i>	<i>Commenters</i>
USA	35.6%
UK	7.5%
Canada	4.9%
Germany	4.8%
The Netherlands	3.7%
Italy	3.7%
Brazil	3.3%
France	2.6%
Mexico	2.4%
Spain	2.4%
Australia	2.2%
Sweden	1.8%
Finland	1.2%
Malaysia	1.1%
Poland	1.0%
Argentina	1.0%
Philippines	1.0%
Romania	1.0%

Individual comments

This subsection gives broad summary statistics about comments, again to serve as context for this study and for future investigations into YouTube commenting.

Length The average length of comments is quite short at 95.5 characters, including spaces and punctuation (see also Figure 2). The most common length is 19 characters and the median is 58 characters (about 11 words). The nominal maximum length for YouTube comments seems to be 500 characters, although a few comments exceeded this (see Figure 3) because the program that downloaded the comments standardised the characters by converting tabs to five spaces and converting line end characters to HTML
 codes. About 95% of comments have lengths less than 344 characters (about 65 words).




Figure 2. Comment lengths for a random comment from each video retrieved (one selected at random per search) with at least one comment.

Sentiment From the SentiStrength results (see the methods section), the apparently English YouTube comments tend to be mildly positive: the mean average strength of positivity on a scale of 1 (no positivity) to 5 (strongly positive) was 2.01 whereas the mean for the equivalent negativity scale was 1.50 (only half as far along the scale). Figure 3 shows that strong sentiment is rare, but suggests that negative strong sentiment is more common than positive strong sentiment, even though most comments contain no negativity.




Figure 3. Sentiment strength for 1,242,885 predominantly English YouTube comments.

All comments for a video

This subsection gives some basic statistics about discussions and the following subsection focuses on more in-depth analyses.

Length The complete comment sets sample was used to analyse the characteristics of complete collections of comments associated with individual videos. The complete comment sets sample contained an average of 76.2 comments per video (the average was 108.9 comments per video for the larger comment sets sample).

Sentiment To examine the role of sentiment in YouTube comments, the average level of positive and negative sentiment strength was calculated for each set of 9,592 commented videos and correlated against the number of comments extracted. Unsurprisingly, average positive and negative sentiment strengths were negatively correlated for videos (Spearman's rho -0.213, $p=0.000$), but this shows that videos either tend to have positive comments or negative comments rather than expressive comments (i.e., high positive or negative sentiment strengths) or neutral comments. The number of comments to a video correlated with average negative sentiment strength (Spearman's rho 0.242, $p=0.000$) and negatively correlated with positive sentiment strength (Spearman's rho -0.113, $p=0.000$). Hence videos with many comments tend to have disproportionately strong negative sentiments expressed in them – probably because a debate is occurring in the comments. In contrast, positive sentiment is disproportionately strong among videos with fewer comments, perhaps suggesting that either positive comments rarely trigger reactions or that viewers feel little need to register positive comments on a video that already has some.

Replies as a proxy for discussions

Reply densities were calculated for the *complete comment sets* sample (all comments for a video with 2-998 comments), with the assumption that the proportion of replies to previous comments is a reasonable indicator of the extent of discussion between commenters on a video. The average reply density (see methods) was 0.234 (i.e., 23.4% of YouTube comments in complete comment sets are replies, if there was a previous comment that they could be a reply to). There is a significant correlation between discussion size (estimated total number of comments for the video, as reported by the YouTube API) and density (Spearman's rho = 0.548, $p = 0.000$). From Figure 4, the reply density is approximately constant at 0.285 for 50-998 comments returned, but increases from approximately 0.15 in a logarithmic curve shape between 2 and 50 comments. People watching a YouTube video

will see about 9 of the most recent comments by default, in addition to perhaps one or two previous highly-rated comments. Hence, it seems unlikely that many viewers would reply to a comment that was older than 9 comments, unless it was highly rated. The curve extends beyond 9 to 50 and this suggests that something inherent in some videos attracts many comments, rather than a natural self-organising feedback process in which the primary driving process is that existing comments attract more comments. Note that the approximate reply density for the excluded videos, calculated using the same formula $[\#replies/(\#comments-1)]$, despite the reservations given above, was slightly higher at 0.265, with the density for the combined set being 0.235.




Figure 4. Average reply density $\#matching_replies/(\#comments_extracted - 1)$ from 35,347 YouTube videos with 2-998 comments. The data is binned into 100s so that each set of five points (except the last, for 80 videos) represents a minimum of 100 videos. Each vertical stack of five points is plotted against the average number of comments for the bin. Bins are chosen to be not overlapping; for example there is one bin for all 2718 videos with 2 comments and one bin for all 103 videos with 883-944 comments. For the latter and to explain the other data in the graph, the minimum density is 0.04, 95% of videos had a density of at least 0.08 (Lower 95%), the average density was 0.29, 95% of videos had a density of less than 0.57, and the maximum density was 0.85. Note that the maximum and minimum values are misleading for videos with under 40 comments (mostly 1s and 0s respectively in the figure) because these are based upon significantly more than 100 videos (177-2718).

Figure 4 reveals that the density of discussions varies significantly, even for videos with many comments, and the spread from the lower 95th percentile to the upper 95th is such that it does not seem to be reasonable to claim that there is a typical density of discussion: A video with a reply density of approximately 5%-55% would seem normal in this respect. Note that for videos with over 998 comments returned by the API (and some with many more than 1,000 comments), the average density of matching replies is slightly lower at 0.265, but the real figure may be higher as some of the first 1000 comments will be replies to earlier comments. At all discussion sizes, some videos have few replies and some have many. More specifically, and as a benchmark, for videos with 50-995 comments: 90% have a reply density between 0.075 and 0.546.

Categories and content In order to gain insights into the types of videos attracting the most and least replies (see the Appendix for examples), the 100 videos with the highest reply density were selected and compared to the 100 videos with the lowest reply density in terms of their official YouTube categories, as listed below each video on its home page (Figure 5). A minimum threshold number of

comments of 250 per video was set to eliminate videos with too few comments to have a reliable idea about the density of discussion generated by them. The results show clear differences in terms of the most common categories. It seems that Music and Comedy videos attract the least replies, as well as How to & Style videos. In contrast, the most discussed topics are News & Politics and Science & Technology. Some of these differences seem logical at face value; for instance, music, comedy and entertainment seem to be passive media consumption activities and so people choosing these options may not wish to engage. In contrast, News & Politics seems to be a natural topic for discussion. The dense discussions for Science & Technology and Education are perhaps more surprising, but 10 out of the 14 dense reply Education videos were about religion (including one about evolution) and one was about politics, so the categorisation was perhaps misleading for this group. Similarly, 8 of the 18 dense reply Science & Technology videos were about religion, evolution or creationism, suggesting that this was a major cause of this category's dense replies. Nevertheless, other dense reply science videos discussed climate change (3), and space or astrophysics (4), indicating that some hard science topics can also attract a significant amount of replies. Many of the other dense reply videos were about religion: 3 in Nonprofits & Activism, 4 in News & Politics, and 5 in People & Blogs, making a total of 30 religion-related videos in this group. No other topic attracted a similar number of videos. The second most popular broad topic in the dense reply group was the economy or the economic crisis, with 5 videos.

There was also a small but significant difference in the popularity of the two different groups of videos: the dense reply set attracted 92.6% "likes" while the sparse reply set attracted 96.0%. Also, the dense set attracted significantly fewer ratings; a median of 252.5 in comparison to 796. This would be consistent with the sparse discussion videos being almost universally popular, at least amongst those who viewed them, and triggering uncontroversial statements of approval amongst the minority of viewers that left a comment. The high approval ratings for the most discussed videos indicates that the discussions may tend to involve a small number of people that disagree with the majority view of those finding the video.




Figure 5. YouTube categories for the 100 videos with the highest/lowest reply densities, (videos with 250-998 comments only).

Categories and content for videos with 999+ comments The above discussion of categories and content excluded videos with 999+ comments because their reply densities could not be accurately calculated without the missing comments. This section discusses the categories and content for videos with 999+ comments using estimated reply densities and comparing the results to those above for videos with 250-998 comments. Figure 6 reports the categories for the 50 videos from the 999+ comments set with the densest discussions and the 50 videos from the 999+ comments set with the least dense discussions. This additional data set was made from 100 rather than 200 videos altogether to give approximately the same range of densities as with the 250-998 comments data set. The results are broadly similar for both data sets except for a few differences. Absent from the 250-998 comments

data set, the Shows category is a significant presence in the 999+ comments data set (30%), attracting mainly low density discussions. This imbalance reflects that of the similar Entertainment category in the 250-998 comments data set, although the Entertainment category is not unbalanced in the 999+ comments data set (the three high reply density Entertainment videos triggered discussions on religion, right wing politics and Windows vs. Linux). Presumably videos in the Shows category are typically popular and attract 999+ comments due to their mass media associations. The other large difference is that the 999+ comments data set has fewer Music category videos and these are evenly spread between the high and low reply density cases (8% in both cases). The four high reply density videos were not discussed for their music or musicians, however: one triggered a political discussion (a song about Yugoslavia), the second was about religion, the third was a death metal song but the comments discussed death metal in general, and the fourth was a comedy song with comments about racism, culture and national differences. This suggests that when music triggers significant debates the causes may not be the music itself. Additional inspection of the 999+ comment videos revealed several low reply density videos (18%) containing competitions requiring the viewer to leave a comment to enter. Religion was well represented in the high reply density videos (36%), as was politics (34%). Science (14%) and climate change (4%) were again represented, but the economy was not. In summary, it would be reasonable to claim that the themes identified for the 250-998 comments data set are broadly consistent with the results from the 999+ comments data set, but with the latter containing many more competitions and Shows videos.




Figure 6. YouTube categories for the videos with the highest/lowest reply densities, broken down by whether the video returned 250-998 comments or 999+ comments in the YouTube API. The main data is for the 250-998 comments set (200 videos in total; copied from Figure 5) and the secondary data is for the 999+ comments set (100 videos in total).

Limitations

A key limitation of this research is that it is not based upon a random sample but on searches from a list of predominantly English terms. This causes biases since the results are impacted by the YouTube ranking algorithm and the word list approach, which causes its own biases. The sections covering videos attracting less than the maximum number of comments automatically accessible via YouTube excluded about 4% of the longest discussions - those with over 998 comments - these were analysed separately for categories and content. Although this is a numerically insignificant number of videos, these long discussions may have unusual characteristics that may not be represented in the remainder of the data.

A more general limitation is that the results are based upon convenience data in the sense that the factors analysed are those that happen to be reported by YouTube (e.g., commenter age, gender and location), ignoring any factors that were not reported but which are nevertheless important (e.g., reason for joining YouTube). In addition, most of the data analysed is self-reported and some is deliberately incorrect.

For the sentiment analysis, a limitation is that the algorithm used for this is imperfect and therefore the sentiment results are not likely to be completely accurate. Nevertheless, the computer program used has about the same level of accuracy as humans (Thelwall et al., 2010) and, unlike most sentiment analysis algorithms, does not pick up topic words but only directly detects expressions of sentiment and therefore should not give systematic biases unless there are videos that attract complex expressions of sentiment that the program cannot detect. This is most likely to be relevant to political discussions, in which sarcasm can be expected.

For the discussion of reply densities, an important limitation is that some users might reply to other comments without using the official reply function. Hence the calculated density of replies might be underestimated in many cases. Related to this, the replies may sometimes be part of a discussion or debate but in other times they might be simple agreements. Although the prevalence of controversial topics, like religion, in the results and the association between negative sentiment and denser discussions suggest that the dense replies are part of a genuine debate, this has not been proven in each case.

The categories and content discussion for high and low reply density videos excludes the 88% of videos with 1-249 comments because a large number of comments is needed to reliably decide whether a discussion is dense or not. For this analysis, the excluded data (videos with 1-249 comments; 88% of the total) account for only 9% of comments to videos, the main analysed data set (videos with 250-998 comments; 8% of the total) accounts for 10% of comments, and the secondary data set (videos with 999+ comments; 4% of the total) accounts for 81% of comments made to videos. Hence, the content findings collectively cover the majority of comments (10% + 81% = 91%) but a minority of commented videos (8% + 4% = 12%).

Conclusions

The investigations of YouTube comments, commenters and discussions have given baseline statistics to aid future readers in assessing the extent to which any videos analysed are typical. From the English-dominated sampling method, YouTube commenters predominantly state a male gender (72.2%) and have a median stated age of 25. YouTube comments are predominantly short, with a median of 58 out of a possible 500 characters (about 11 words). This suggests that comments are deliberately kept short rather than being constrained to be short (probably in contrast to Twitter). Typical comments are mildly to moderately positive, although 35% of comments contain some negativity. Videos attracting text responses (but less than 999) had an average of 76.2 comments. Although negative sentiment was uncommon, it was more prevalent in comments for videos attracting many comments; conversely positive sentiment was disproportionately common in videos attracting few comments. Thus, it seems that negativity can drive commenting – perhaps partly through long-running acrimonious comment-based discussions.

In terms of the density of replies to comments on a video, there was a wide variety and 90% of discussion densities varied between 0.075 and 0.546. This confirms the heterogeneity of YouTube, but means that researchers investigating videos in the future would need to find a discussion density of over 0.546 to prove statistically that they had attracted unusually dense discussions. Although about a quarter of comments attracted a reply, this fraction varied greatly by video, with many videos having few replies to comments and many videos having replies to a majority of comments. It seems that the topic of a video is a key determinant of whether it will create much discussion in the sense of a high proportion of comments being replies to previous comments. Amongst videos attracting 250-998 comments (i.e., the range for which extreme reply density videos could be reliably determined), the single topic attracting the highest proportion of replies per comment was religion, accounting for 30% of the 100 videos with most replies per comment. In contrast, music and comedy videos together accounted for the majority of videos attracting few replies per comment. Nevertheless, a range of other topics also attracted many replies per comment, particularly within the broad categories of News

& Politics, and Science & Technology. This would be consistent with the hypothesis that there are different audiences for YouTube: some come to be passively entertained and don't engage significantly with other users, whereas others are prepared to engage in discussion around controversial or interesting topics. The same seems to be true for videos attracting 999+ comments (videos attracting under 250 comments could not be easily analysed). This aligns with claims that audiences consume media in different ways to support their own personal goals (Blumler & Katz, 1974).

The extent of interaction between YouTube commenters is remarkable: just under a quarter of comments on a video after the first were replies to previous comments. This suggests that YouTube hosts genuine audience discussions about the various topics hosted on the site. As the examples in the appendix show, some of these are genuine debates on controversial issues, which raises the possibility that YouTube is a significant public space (or even a public sphere, Habermas, 1991) for engaging in debate and exchanging opinions. The high popularity of YouTube and the finding that far more people discuss videos offline than comment on them online for some topics (Milliken, Gibson, O'Donnell, & Singer, 2008; Milliken, Gibson, & O'Donnell, 2008) suggests that such discussions may be socially significant even though under 0.5% of viewers leave a comment. Additional research is needed to investigate this issue for different discussion topics within YouTube. Moreover, the nature of debates that occur in YouTube is unclear. For example, it is awkward and takes time for a user to access all the comments on a YouTube video if there are more than about 10 and they have to be paged through on the site. Hence, it seems highly unlikely that a popular video would host a single coherent debate but it may be possible for videos to host numerous debates between small groups of commenters. Perhaps such debates would only be possible in real time for the most popular videos because it may be too difficult for a user to find replies to their comments otherwise.

The findings summarised here fulfil the goal of the paper to set benchmarks against which future qualitative or quantitative research can be checked. In particular, those investigating a video can use the reply density formula to see whether the comments on it form an unusually dense discussion or not, or could use SentiStrength to assess whether the sentiment content of the comments is similar to the rest of YouTube. An important additional implication of the findings for future YouTube research is that the site should not be treated as an undifferentiated mass but as a place that is used by different audiences in different ways. In particular, when analysing a particular video or set of videos it would be best to benchmark it or them against videos from the same genre rather than against a random sample of YouTube videos; this would give a better idea of any unusual features.

Acknowledgements

This work was supported by a European Union grant by the 7th Framework Programme, Theme 3: Science of complex systems for socially intelligent ICT. It is part of the CyberEmotions project (contract 231323).

Appendix – Examples of extreme discussion density videos

An example of a video with an unusually high discussion density is⁶ zdFVAUCM6X4, "Skeptics Among Us: Atheists Visit The Creation Museum - Part 1 of 3". The visit was designed to trigger discussions about religion and achieved this with a density of 0.64 from the 993 comments. Another video is B18-YC8oPiE, "El Big Bang, El tiempo, y el Creador (1 de 2)", which has a density of 0.58 and features a discussion between creationism and atheism in English with Spanish subtitles. A third is 6b2gswxOomQ, "Dialog Pindah Kuil Kecoh: Khalid diboo!!" in Malay, with a density of 0.67, featuring a news story discussing a contentious plan to set up Hindu temples in a particular area of Malaysia.

An example of a video with a low discussion density is pf4hcAhIDjU, "Erkin Koray - Öyle Bir Geçer Zaman Ki", from a Turkish rock singer with a career spanning 50 years. Comments tended to be simple messages of appreciation (in Turkish), such as "I love this guy's songs". Another example is the comedy video KWEbRNwvJT's "Ventrilo Rapage - Vent Virus", which attracted mainly positive comments such as, "Hilarious man". Finally, another music video bh9XefYUgoc, "ricardo

⁶ Add the YouTube ID to the end of the base URL <http://www.youtube.com/watch?v=> to access the video

arjona-te conozco", attracted mainly positive comments (in Spanish) like "this song was successful in its time and today it is an excellent classic". It is by Grammy award-winning Guatemalan singer Ricardo Arjona, who first became popular in 1989.

References

- Alexander, J., & Losh, E. (2010). "A YouTube of one's own?": "Coming out" videos as rhetorical action. In C. Pullen & M. Cooper (Eds.), *LGBT Identity and online new media* (pp. 37-50). New York, NY: Routledge.
- Alonzo, M., & Aiken, M. (2004). Flaming in electronic communication. *Decision Support Systems*, 36(3), 205-213.
- Ayres, C. (2009). Revenge is best served cold on YouTube. *The Times Online*(July 22, 2009), Retrieved September 1, 2009 from: http://www.timesonline.co.uk/tol/comment/columnists/chris_ayres/article6722407.ece.
- Baumgartner, J. C., & Morris, J. S. (2010). MyFaceTube politics: Social networking web sites and political engagement of young adults. *Social Science Computer Review*, 28(1), 24-44.
- Blumler, J. G., & Katz, E. (1974). *The uses of mass communications: Current perspectives on gratifications research*. Beverly Hills, CA: Sage.
- Burgess, J., & Green, J. (2009). *YouTube: Online video and participatory culture*. Cambridge: Polity.
- Castells, M. (2008). The new public sphere: Global civil society, communication networks, and global governance. *The ANNALS of the American Academy of Political and Social Science*, 616(1), 78-93.
- Cha, M., Kwak, H., Rodriguez, P., Ahn, Y.-Y., & Moon, S. (2007). I tube, you tube, everybody tubes: Analyzing the world's largest user generated content video system. *Internet Measurement Conference 2007*, Retrieved September 23, 2009 from: <http://www.imconf.net/imc-2007/papers/imc2131.pdf>.
- Cha, M., Kwak, H., Rodriguez, P., Ahn, Y.-Y., & Moon, S. (2009). Analyzing the video popularity characteristics of large-scale user generated content systems. *IEEE/ACM Transactions on Networking*, 17(5), 1357-1370.
- Cheng, X., Liu, J., & Dale, C. (in press). Understanding the characteristics of internet short video sharing: A YouTube-based measurement study. *IEEE Transactions on Multimedia*.
- Chmiel, A., Sienkiewicz, J., Paltoglou, G., Buckley, K., Thelwall, M., & Holyst, J. A. (2011). Negative emotions boost user activity at BBC forum. *Physica A*, 390(16), 2936-2944.
- Cornelius, R. R. (1996). *The science of emotion*. Upper Saddle River, NJ: Prentice Hall.
- Derks, D., Bos, A. E. R., & von Grumbkow, J. (2008). Emoticons and online message interpretation. *Social Science Computer Review*, 26(3), 379-388.
- Ding, Y., Jacob, E. K., Zhang, Z., Foo, S., Yan, E., George, N. L., et al. (2009). Perspectives on social tagging. *Journal of the American Society for Information Science & Technology*, 60(12), 2388-2401.
- Fauconnier, S. (2011). Video art distribution in the era of online video. In G. Lovink & R. Somers Miles (Eds.), *Video Vortex Reader II* (pp. 108-125). Amsterdam: Institute of Network Cultures.
- Friedman, B., Khan, P. H., & Howe, D. C. (2000). Trust online. *Communications of the ACM*, 43(12), 34-40.
- Gill, P., Arlitt, M., Li, Z., & Mahanti, A. (2007). YouTube traffic characterization: A view from the edge. *Internet Measurement Conference 2007*, Retrieved September 23, 2009 from: <http://www.imconf.net/imc-2007/papers/imc2078.pdf>.
- Gueorguieva, V. (2008). Voters, MySpace, and YouTube: The impact of alternative communication channels on the 2006 election cycle and beyond. *Social Science Computer Review*, 26(3), 288-300.
- Habermas, J. (1991). The public sphere. In C. Mukerji & M. Schudson (Eds.), *Rethinking popular culture: contemporary perspectives in cultural studies* (pp. 398-404). Berkeley: Univ. of California Press.
- Hang, W.-y., & Yun, S.-y. (2008). How User-Generated Content (UGC) campaign changes electoral politics? *Korea Observer*, 39(Autumn), 369-406.

- Herring, S. C. (2002). Computer-mediated communication on the Internet. *Annual Review of Information Science and Technology*, 36, 109-168.
- Hopkins, J. (2006). Surprise! There's a third YouTube co-founder. USA Today, Retrieved July 15, 2011 from: http://www.usatoday.com/tech/news/2006-2010-2011-youtube-karim_x.htm.
- Huberman, B. A., Romero, D. M., & Wu, F. (2009). Crowdsourcing, attention and productivity. *Journal of Information Science*, 35(6), 758-765.
- Im, H.-B. (2010). Development and change in Korean democracy since the democratic transition in 1987. In Y.-w. Chu & S.-l. Wong (Eds.), *East Asia's new democracies: deepening, reversal, non-liberal alternatives* (pp. 102-121). London: Routledge.
- Java, A., Song, X., Finin, T., & Tseng, B. (2007). Why we twitter: understanding microblogging usage and communities. In *Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining and social network analysis* (pp. 56-65). New York, NY: ACM Press.
- Kwak, H., Lee, C., Park, H., & Moon, S. (2010). What is Twitter, a social network or a news media? In *Proceedings of the 19th international conference on world wide web* (pp. 591-600). New York, NY: ACM Press.
- Lange, P. G. (2007a). Commenting on comments: Investigating responses to antagonism on YouTube. *Annual Conference of the Society for Applied Anthropology*, Retrieved April 6, 2011 from <http://sfaapodcasts.files.wordpress.com/2007/2004/update-apr-2017-lange-sfaa-paper-2007.pdf>.
- Lange, P. G. (2007b). Publicly private and privately public: Social networking on YouTube. *Journal of Computer-Mediated Communication*, 13(1), Retrieved May 8, 2008 from: <http://jcmc.indiana.edu/vol2013/issue2001/lange.html>.
- Lange, P. G. (2008). Living in YouTubia: Bordering on civility. *Proceedings of the Southwestern Anthropological Association Conference*, 98-106.
- Lange, P. G. (2009). Videos of affinity on YouTube. In P. Snickars & P. Vonderau (Eds.), *The YouTube Reader* (pp. 228-247). Stockholm: National Library of Sweden.
- Lazzara, D. L. (2010). YouTube courtship: The private ins and public outs of Chris and Nickas. In C. Pullen & M. Cooper (Eds.), *LGBT Identity and online new media* (pp. 51-61). New York, NY: Routledge.
- Lewis, S. P., Heath, N. L., St Denis, J. M., & Noble, R. (2011). The scope of nonsuicidal self-injury on YouTube. *Pediatrics*, 127(3), e552 -e557.
- Lo, A. S., Esser, M. J., & Gordon, K. E. (2010). YouTube: a gauge of public perception and awareness surrounding epilepsy. *Epilepsy & Behavior*, 17(4), 541-545.
- Losh, E. (2008). Government YouTube bureaucracy, surveillance, and legalism in state-sanctioned online video channels. In G. Lovink & S. Niederer (Eds.), *Video Vortex Reader* (pp. 111-124). Amsterdam: Institute of Network Cultures.
- Madden, M. (2007). Online video. *Pew Internet*, Retrieved June 17, 2011 from: <http://www.pewinternet.org/Reports/2007/Online-Video.aspx>.
- Maia, M., Almeida, J., & Almeida, V. (2008). Identifying user behavior in online social networks. *Proceedings of the 1st Workshop on Social Network Systems*, 1-6.
- Milliken, M., Gibson, K., O'Donnell, S., & Singer, J. (2008). User-generated online video and the Atlantic Canadian public sphere: A YouTube study. In *Proceedings of the International Communication Association Annual Conference*. Retrieved July 25, 2011 from: <http://nparc.cisti-icist.nrc-cnrc.gc.ca/npsi/ctrl?action=rtdoc&an=8913990&lang=en>.
- Milliken, M., Gibson, K., & O'Donnell, S. (2008). User-generated video and the online public sphere: Will YouTube facilitate digital freedom of expression in Atlantic Canada? *American Communication Journal*, 10(3), Retrieved July 25, 2011 from: http://ac-journal.org/journal/pubs/2008/Fall%2008%2020-%2020Defining%2020Digital%2020Freedom/Article_2015.pdf.
- Mislove, A., Marcon, M., Gummadi, K. P., Druschel, P., & Bhattacharjee, B. (2007). Measurement and analysis of online social networks. *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement* 29-42.
- Molyneaux, H., O'Donnell, S., Gibson, K., & Singer, J. (2008). Exploring the gender divide on YouTube: An analysis of the creation and reception of vlogs. *American Communication*

- Journal*, 10(2), Retrieved March 1, 2011 from: <http://iitatlns2012.iit.nrc.ca/iit-publications-iti/docs/NRC-50360.pdf>.
- Naveed, N., Gottron, T., Kunegis, J., & Alhadi, A. C. (2011). Bad news travel fast: A content-based analysis of interestingness on Twitter. *WebSci 2011*, Retrieved July 16, 2011 from: http://www.websci2011.org/fileadmin/websci/Papers/2050_paper.pdf.
- Paolillo, J. C. (2008). Structure and network in the YouTube core. *Proceedings of the 41st Annual Hawaii International Conference on System Sciences*, Retrieved June 16, 2011 from: <http://www.computer.org/portal/web/csdl/doi/2010.1109/HICSS.2008.2415>
- Purcell, K. (2010). The state of online video. *Pew Internet*, Retrieved June 15, 2011 from: <http://www.pewinternet.org/Reports/2010/State-of-Online-Video.aspx>.
- Shifman, L. (in press). An anatomy of a YouTube meme. *New Media and Society*.
- Siersdorfer, S., Chelaru, S., Nejd, W., & Pedro, J. S. (2010). How useful are your comments?: Analyzing and predicting youtube comments and comment ratings. *Proceedings of the 17th international conference on World Wide Web*, Retrieved June 16, 2011 from: <http://www.l2013s.de/~siersdorfer/sources/2010/wfp0542-siersdorfer.pdf>.
- Sobkowicz, P., & Sobkowicz, A. (2010). Dynamics of hate based Internet user networks. *European Physics Journal B*, 73(4), 633-643.
- Sommerville, Q. (2009). China 'blocks YouTube video site. BBC News, Retrieved July 15, 2011 from: <http://news.bbc.co.uk/2012/hi/asia-pacific/7961069.stm>.
- Stauff, M. (2009). Sports on YouTube. In P. Snickars & P. Vonderau (Eds.), *The YouTube Reader* (pp. 236–251). Stockholm: National Library of Sweden.
- Steinberg, P. L., Wason, S., Stern, J. M., Deters, L., Kowal, B., & Seigne, J. (2010). YouTube as Source of Prostate Cancer Information. *Urology*, 75(3), 619-622.
- Sunstein, C. R. (2007). *Republic.com 2.0*. Princeton: Princeton University Press.
- Sykora, M. D., & Panek, M. (2009a). Media sharing websites and the US financial markets. IADIS International Conference WWW/Internet, Retrieved July 18, 2011 from: <https://dspace.lboro.ac.uk/dspace-jspui/handle/2134/6423>.
- Sykora, M. D., & Panek, M. (2009b). Financial news content publishing on Youtube.com. Proceedings of the 3rd International Workshop on Soft Computing Applications, 99-104, Retrieved July 118, 2011 from: <https://dspace.lboro.ac.uk/dspace-jspui/handle/2134/6420>.
- Thelwall, M., Buckley, K., Paltoglou, G., Cai, D., & Kappas, A. (2010). Sentiment strength detection in short informal text. *Journal of the American Society for Information Science and Technology*, 61(12), 2544-2558.
- Thelwall, M., & Wilkinson, D. (2010). Public dialogs in social network sites: What is their purpose? *Journal of the American Society for Information Science & Technology*, 61(2), 392-404
- Thelwall, M., Wouters, P., & Fry, J. (2008). Information-centred research for large-scale analysis of new information sources. *Journal of the American Society for Information Science and Technology*, 59(9), 1523-1527.
- Thorson, K., Ekdale, B., Borah, P., Namkoong, K., & Shah, C. (2010). YouTube and Proposition 8: A case study in video activism. *Information, Communication & Society*, 13(3), 325-349.
- Tremayne, M., Zheng, N., Lee, J. K., & Jeong, J. (2006). Issue publics on the web: Applying network theory to the war blogosphere. *Journal of Computer-Mediated Communication*, 12(1), 290-310.
- Van Langendonck, G. (2009). Iconic Iran video was posted in the Netherlands. *NRC Handelsblad*, Retrieved April 13, 2011 from: http://vorige.nrc.nl/international/article2280315.ece/Iconic_Iran_video_was_posted_in_the_Netherlands.
- van Zoonen, L., Mihelj, S., & Vis, F. (in press). YouTube interactions between agonism, antagonism and dialogue: Video responses to the anti-Islam film Fitna. *New Media and Society*.
- van Zoonen, L., Vis, F., & Mihelj, S. (2010). Performing citizenship on YouTube: Activism, satire and online debate around the anti-Islam video Fitna. *Critical Discourse Studies*, 7(4), 249-262.
- Waldfoegel, J. (2009). Lost on the web: Does web distribution stimulate or depress television viewing? *Information Economics and Policy*, 21(2), 158-168.

- Walther, J., & Parks, M. (2002). Cues filtered out, cues filtered in: computer-mediated communication and relationships. In M. Knapp, J. Daly & G. Miller (Eds.), *The Handbook of Interpersonal Communication (3rd ed.)* (pp. 529-563). Thousand Oaks, CA: Sage.
- Wasserman, S., & Faust, K. (1994). *Social network analysis: Methods and applications*. Cambridge, NY: Cambridge University Press.