Comment Analysis of YouTube videos Discussing Deliberate Self-harm

Muhammad Abubakar Alhassan[1] [0000−0002−7497−6941] and Diane Pennington[2] [0000−0003−1275−7054]

University of Strathclyde, Department of Computer and Information Sciences, 26 Richmond St, Glasgow G1 1XH, United Kingdom.
muhammad.alhassan@strath.ac.uk diane.pennington@strath.ac.uk

Abstract. Online social media platforms remain an excellent source of data for information scientists. Existing studies have found that people who self-harm find it easier to disclose information regarding their behaviour on social media as compared to in-person interactions. Due to the large and growing volume of user-generated content on YouTube, sources of videos presenting information concerning self-harm and discussions surrounding those videos could be hidden by other contents. By using a categorisation codebook and state-of-the-art topic and sentiment analysis techniques, the authors identified distinct groups of users who uploaded videos about self-harm on YouTube (n=107) and uncovered the topics and sentiments expressed in 27,520 comments. In addition to other sources, our investigations discovered that 56% of the people uploading the examined videos are *non-professionals*, in contrast to the group of *professionals* with only 11% of the videos in the sample. In grouping comments based on similar topics, we discovered that *self-harming users*, *clean (recovered) users*, *at-risk audiences*, and *appreciative users* responded to the examined videos. Viewers responded more positively to '*recovered from self-harm*' and '*appreciative*' responses, as opposed to '*at-risk*' and '*self-harm*' comments with a high negative sentiments. These features could be used to build a classifier, although more research is needed to investigate self-injurious information to better support digital interventions for effective prevention and recovery.

Keywords: social media · YouTube · self-harm · self-injury

1 Introduction

The ubiquitous nature of social media may create opportunities that could help in mitigating undiagnosed mental health problems. Many studies have explored social media use and its factors associated with mental well-being. Most focus on anxiety, depression and eating disorders [5, 7], with less attention on self-harming behaviours [6,25]. In the last decade, as a response to a rising number of case reports by physicians, schools and communities, investigating self-harming behaviours has gained more attention from researchers [23]. Self-harm is defined as intentional physical harm to oneself with no intent to end their life. This behaviour is most common among young people, and it usually involves cutting, carving, or burning of the skin [23]. Also, self-mutilation such as headbanging and bruising are unacceptable [11].

In this study, we applied the definition of self-harming behaviour as reported by the Child and Adolescents Self-harm in Europe (CASE) research [16]. The study established criteria for identifying individuals who deliberately self-harm with a self-reported questionnaire to participants in six European countries. The authors constructed a standard definition that characterised deliberate self-harming (DSH) behaviours, including cutting oneself or ingesting dangerous substances, and developed a concrete approach for investigating the prevalence of self-harm. Many researchers use this definition as it entails most common forms of intentional self-injury [1,4]. Nevertheless, the stereotype and stigma surrounding DSH can make the Internet a stimulating space for those engaged in it. The anonymity and boundary-spanning of online networks can allow people greater freedom to convey views deemed hard to share [18]. In other words, online social networks connect people with commonalities much more than in offline social interactions [20, 27]. Young people who self-harm are even more likely to participate in online activities than those who do not [22].

Some YouTube users record and share videos about their personal experiences or stories. YouTube is among the online platforms on which self-harming individuals share experiences, such as normalising self-harm (*dangerous*) or exchanging peer support (*helpful*) [12]. Further research is needed to understand the various sources of videos about self-harm on YouTube, as well as what viewers discuss and their opinions about the videos expressed through their comments. In line with this, our study aimed to answer the following research questions:

- *Who uploads videos discussing self-harm on YouTube, and how do viewers rate the videos?*
  With this question, we aimed to explore the groups of people involved in disseminating videos about self-harm on YouTube. Through exploring the characteristics of their videos, we wanted to understand how their audiences rated the videos on various channels.
- *What are viewers discussing concerning videos presenting self-harm information on YouTube?*
  Here, we sought to uncover hidden topics viewers discussed, and we grouped comments based on similar topics. Addressing this question increased our understanding of the conversations surrounding videos related to self-harm on YouTube.
- *What opinions do viewers communicate in their comments?* To answer this question, viewer's sentiments from the identified topics were examined. The opinions of the audiences viewing videos about self-harm is essential to explore. This will help to understand *dangerous* or *helpful* opinions to promote or prevent self-harm on YouTube.

However, this research paper is organised as follows; the next Section highlights the existing studies in this area, and section 3 describes the approach adopted in the study. Section 4 explains our research findings, and the study's discussions and limitations are discussed in Sections 5 and 6. The research conclusions and future directions are presented in the last section of this paper.

2 Deliberate self-harm (DSH) and social media

An increasing amount of research reports that young people who self-harm frequently utilise the internet, with social media specifically being a preferred online communication mode [21]. Recent times have witnessed a growing body of research focusing around the impacts of social media on DSH, increasingly around user-generated content promoting DSH [17].

Many social media users who engage in DSH share content related to their self-harming behaviours, most of which is insufficiently explored through research [28]. One study investigated the differences in characteristics between self-harm and non-self-harm contents on Flickr [28]. Through analysing these and other features, a framework that can automatically detect self-harm content was proposed. Another comparative study investigated self-harm content from Tumblr, Twitter, and Instagram [21]. Using the '#cutting' hashtag, the researchers retrieved self-harm content and analysed 770 posts (333 from Tumblr, 78 from Twitter and 359 from Instagram). Instagram had more graphical content and more users with low self-esteem than Twitter. Furthermore, Twitter was found to have more DSH recovery-related posts than Instagram and Tumblr.

The qualitative study conducted by [8] uncovered attitudes and beliefs of DSH users on Twitter. The author demonstrated that self-harm is not treated lightly on Twitter, and also found posts that escalate the importance of celebrities' behaviour may influence followers' self-harming behaviours. Additionally, the study discovered that social media informs DSH personal stories through videos. In order to determine the potential risks associated with non-suicidal self-injury (NSSI) videos on YouTube, the work of [12] examined the comments made on those videos. Using coding rubrics, the researchers examined comments of the 100 most popular NSSI videos publicly available on YouTube. The authors discovered that the comments on the examined videos were predominantly on self-disclosure, and feedback was directed to the channel or the video uploader. However, the qualitative approach used by the investigators cannot be applied to a large volume of comments, and therefore likely missed crucial information within the comments.

Further effort is needed to understand viewers' opinions and discussions surrounding YouTube DSH videos on a larger scale [8, 28]. This research investigation adds to the current research on the influence of social media in mediating or promoting self-injurious behaviour. We explored various sources of videos concerning self-harm on YouTube. While most of the examined videos were from non-professional YouTubers, we discovered that support organisations, news media agencies, government and non-governmental organisations, and medical and academic experts have contributed to sharing videos about self-harm on YouTube. Our research findings contribute to the body of knowledge on the views and opinions of people commenting on videos concerning intentional self-harm.

3 Methods

We adopted a mixed methods approach for this study. Choosing this depended mainly on the nature of our research questions and the type of data under investigation. In addressing the first research question, this study considered specific criteria explained in Table 1. Following a previous study on YouTube video authorship [19], we started with the author's approach, made a few changes, and then categorised the diverse group of channels included in our sample. We assessed viewers' interactions from each channel. We addressed research questions two and three by utilising a state-of-the-art natural language processing toolkit, a topic modelling algorithm, and a sentiment analysis technique.

Our research utilised the Natural Language Toolkit (NLTK).[1] It is a free and open source package available in Python that incorporates numerous tools for developing a programme and performing data classification tasks. Also, we used the Latent Dirichlet Allocation (LDA) technique, which is a popular probabilistic topic model that identifies similarities linking various data parts [2]. A specific part of the data in topic modelling is considered a textual document, which in our case is an *individual comment* on videos discussing self-harm. Similarly, we performed sentiment analysis of the video responses using the Valence Aware Sentiment Reasoner (VADER), which is a rule-based technique for conducting sentiment analysis of social media textual data [9].

Table 1: Source channel classification scheme

Category	**Description**	**Code**
Professionals	Professional individual(s) appeared in the video content and shared the video on their YouTube channels. Also, their channel's description may have contained information representing medical or academic experts such as psychologists and psychiatrists.	001
Non-professionals	Channel(s) describing the YouTuber as a layperson who promotes mental health awareness, with the video presenter being a non-academic or non-medical expert.	002

[1] https://www.nltk.org/

News media	This category includes channels maintained by traditional news media firms, such as local and international news organisations that have uploaded videos about self-harm.	003
Government organisations	In this group, we included YouTube channels representing government institutions, such as educational or medical institutions.	004
Private organisations	This group is mainly for video channels managed by private companies to promote mental health products, especially directed toward self-harming people.	005
Support organisations	These channels belonged to support organisations, such as the UK-based Samaritans, which have produced videos concerning self-harm to increase awareness, promote support, and encourage recovery.	006

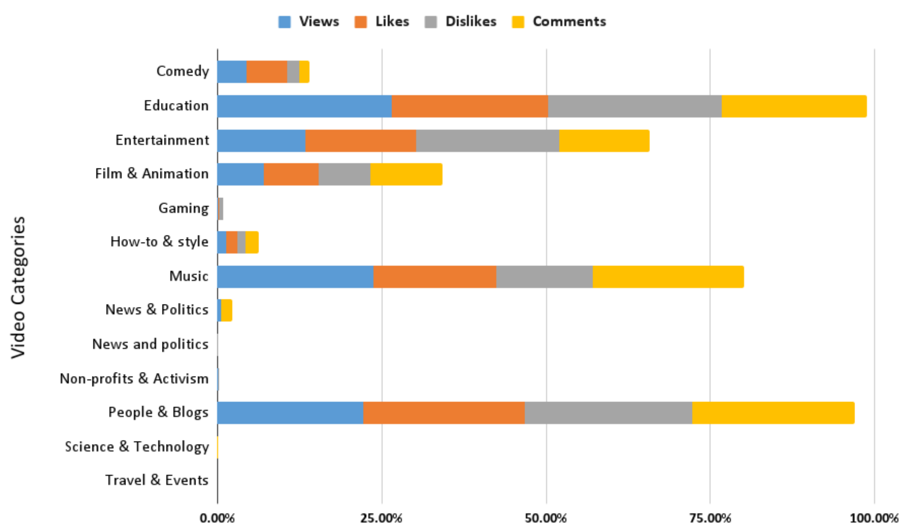

Fig. 1. Percentages of video categories.

3.1 Data collection

As noted above, our approach was informed by previous studies that have investigated YouTube's DSH-related content. YouTube provides access for researchers to extract data via its Application Programming Interface (API). Previous studies used keywords such as 'self-cutting' and 'self-harm' to extract relevant data from YouTube. Our study searched the platform using (1) *'self-harm',* (2) *'self-injury',* (3) *'self-cutting',* (4) *'non-suicidal self-injury',* and (5) *deliberate self-harm'*. Per YouTube's data extraction limits, these terms returned 250 videos (with 50 from each query term) discussing self-harm. Duplicates and non-English videos were removed. This resulted in 172 unique videos uploaded between 2010 and 2020. We extracted a total of 37,100 comments from those videos, and removed duplicate comments or responses written in a language other than English. This produced a total of 27,520 responses from 27,510 unique users for analysis.

3.2 Data analysis

Each of the videos examined featured one or more presenters. Similar to earlier research that demonstrated gender imbalance, showing how females self-harm at a higher rate than males [3, 24], our analysis found that female presenters appeared in 66.67% of the analysed videos, compared to male presenters who appeared in 24.56% of them. As we know, YouTube provides different video categories in which uploaders select the appropriate category related to their content. For example, a tutorial teaching people how to perform mathematical computations would likely be assigned to the *education* category. Therefore, our data analysis explored the categories assigned by the uploader to each of the videos. While video categories could be used to assist users to search specific videos, our analysis identified the most popular categories, including the number of views, likes, and dislikes for each category.

The idea was to call to the attention of the information science community, especially information retrieval system designers and researchers, the utilisation of these attributes to increase the ranking of relevant *helpful* videos for target audiences to increase self-harm awareness and support recovery. We found *people and blogs, education* and *music* to be the most frequently chosen channel categories in our sample. Videos under *education* received fewer likes than those assigned under the *music* category, suggesting that viewers favour entertaining videos over educational videos. Regardless of the channel category, there was a high rate of engagement from viewers, as the videos in our sample received many comments.

3.3 Basic text processing

One of the critical steps of any natural language processing task is basic text processing. This study used python's natural language toolkit and imported all the relevant libraries needed to clean and prepare the set comments for topic modelling and sentiment analysis. In our analysis, we removed stopwords and created a document term matrix from the comments corpus. While this matrix was created using tokenization, it summarises the frequency of each term in every document. We also performed stemming, and lemmatization as the duo are crucial parts of every text pre-processing. Following that, we lemmatized each specific word to its base form, excluding nouns, adjectives, verbs, and adverbs. We retain these parts of speech tags since they are essential in understanding sentence meaning and context. Moreover, to effectively build the topic model, we created a dictionary and corpus.

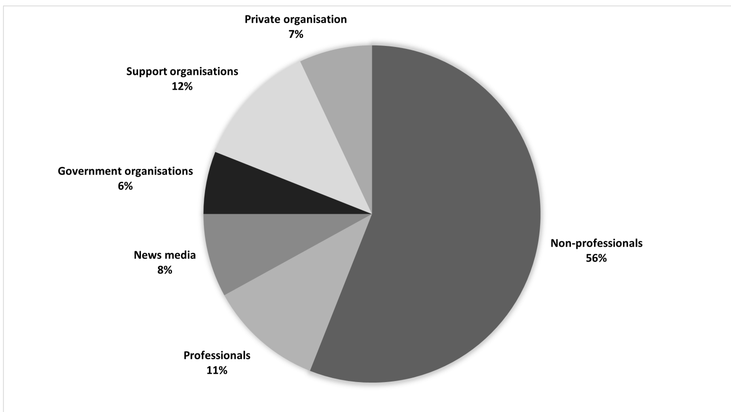

Fig. 2. Video authors' categories

4 Results

4.1 Users sharing videos about self-harm

The present study classified YouTubers who shared videos concerning DSH using the criteria given in Table 1. The proportion of each category is depicted in Figure 2. Non-professional individuals uploaded around 56% of the studied videos, compared to professionals, who authored only 11%. Support organisations accounted for up to 12%, whereas only 6% were from government organisations. Similarly, videos produced by news media organisations and commercial enterprises account for 6% and 7%, respectively.

YouTube viewers engage with the platform in different ways. Firstly, in addition to subscribing to a video channel, there are options for viewers to rate a video through likes and dislikes. Secondly, viewers can interact by commenting on a video and liking or disliking a comment. Our study examined these interactions (except likes and dislikes of comments) from different channel cohorts, as shown in Figure 3. Although all the categories of YouTubers had significant

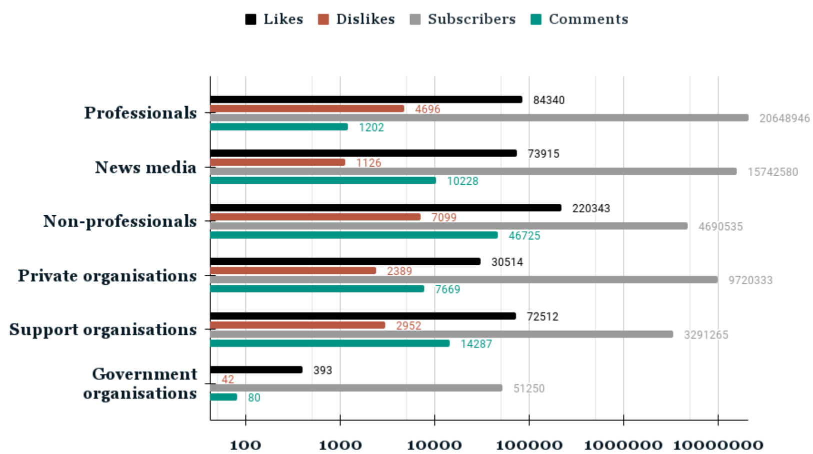

Fig. 3. Groups of users and the frequencies of interactions they received

numbers of subscribers, it is apparent that the audiences favoured (liked) *non-professionals'* video content most.

Similarly, videos from these categories (except those from government institutions) encouraged interactions among viewers, and they received a large number of comments. In contrast to the rest of the categories, the rate of dislikes outweighed comments in the *professionals* group. This suggests that viewers possibly found videos from this group less engaging. Moreover, our analysis explored different video categories assigned by the group of video authors. The matrix table in Figure 4 illustrates the percentage of video categories across the various uploaders' channels. Several channels assigned the *people and blogs* category to one or more videos.

Nevertheless, about 55.2% of the videos uploaded by the group of *non-professionals* were found in the *people and blogs* category. This implies that community members (especially those who had experienced self-harm) published videos to share their experiences and increase awareness about the behaviour. Similarly, videos uploaded in the *education* category appeared in all groups except *news media*, and up to 75% of the videos from *government organisations* were found in the *education* category. Meanwhile, this percentage reduced slightly to 64.3% and 62.5% in the *professionals* and the *support organisations* channels respectively. This shows that these groups are socially active on YouTube through educating the community about self-harm and raising awareness regarding its potential consequences. Additionally, we found no video in the *entertainment* category that was from government organisations or support organisations. Also, *news media* channels comprised the only group containing

	People & Blogs	Education	Entertainment	Film & Animation	Nonprofits & Activism	News & Politics	Music	Howto & Style	Science & Technology	Comedy	Gaming
Non-professionals	**55.2**	11.9	11.9	9.0			4.5	4.5		3	
Professionals	7.1	64.3	**21.4**	7.1							
News media	23.1		**38.5**		7.7	30.8					
Government organisations	12.5	**75.0**			12.5						
Support organisations	25.0	**62.5**			12.5						
Private organisation	**33.3**	16.7	16.7	16.7					16.7		

Fig. 4. Video uploaders and their channel categories

videos from the *news and politics* category, accounting for 30.8%. This percentage increased to 38% in the *entertainment* category. Surprisingly, around 38.5% of videos from the *news media* channels were assigned to the *entertainment* category.

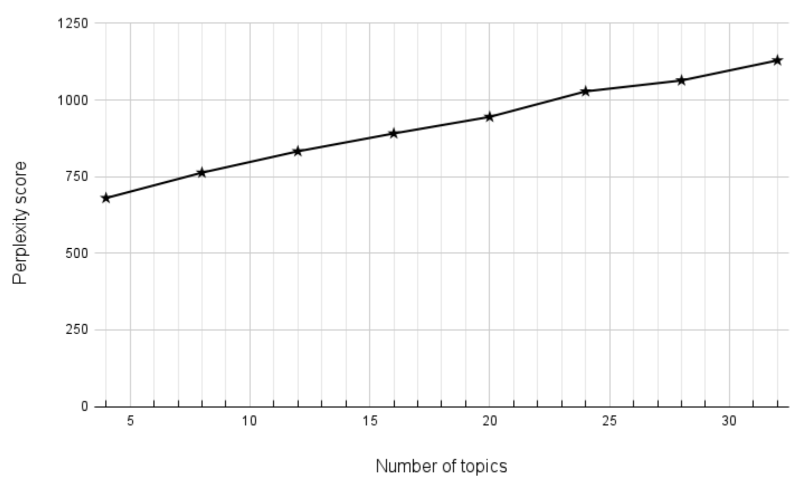

Fig. 5. Choosing appropriate number of topics

4.2 Themes surrounding DSH video responses

To develop an LDA model, a sufficiently large dataset is essential. The minimum size required is determined by the document's attributes and average length - the more significant the dataset, the more accurate the results due to the more significant number of observations. Unlike tweets, YouTube video comments can have a large number of characters. Comment's length could reach up to 10,000 characters. However, our analysis examined 7,520 comments extracted from videos presenting self-harm information on YouTube.

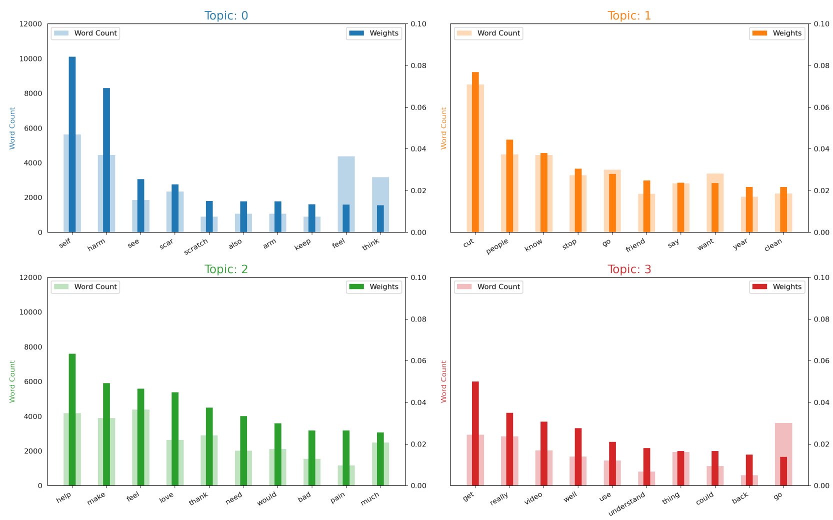

Fig. 6. Terms frequency and their weights

However, although Latent Dirichlet Allocation (LDA) is one of the state-of-the-art techniques for topic analysis [10], determining the correct number of topics in which the LDA model predicts the topic is very challenging. To overcome this issue, we used *perplexity* and ran the model on different topics (4, 8, 16, 20, 24, 28, and 32). Intuitively, the lower the perplexity, the better prediction [2]. After completing different rounds of experiments, a model with four topics produced the best prediction as shown in Figure 4.1. Therefore, we applied this model and analysed the comments corpus, and identified essential topics surrounding videos presenting self-harm information. When interpreting topics, the keywords (unigrams) contained in the topics and their relevance (weights) is significant. The frequency in which the terms appeared in the comments is essential to explore. Intuitively, words that appeared in several topics, including those with a frequency higher than their weight, were considered less relevant as shown in Figure 6. This study considered the first ten terms based on their weights, and interpreted the topics discussed by the viewers as presented in Table 2.

Table 2: Topic interpretations and examples of comments

Topic	Unigrams	Example comments
Topic 0: Self-harming users	self, harm, see, scratch, also, arm, keep, feel, think	1. *"I'm 11 and so far I have cut all my arms, legs and the side of my throat! I'm not proud the scars and blood just show I'm still living and no one is right about me!"* 2. *"I self harm and I often think of ways how to kill myself and how I can self harm even more. It's horrible but I have an addictive personality. My boyfriend self harms as well, we self harm together most the time."*
Topic 1: Clean (recovered) viewers	cut, people, know, stop, go, friend, say, want, year, clean	1. *"I think I'm nearly a year clean from self harm and the kids at my school make fun of self harm all the time and it gets really hard sometimes not to relapse"* 2. *"I think I'm nearly a year clean from self harm and the kids at my school make fun of self harm all the time and it gets really hard sometimes not to relapse"*
Topic 2: At-risk audiences	help, make, feel, thank, need, would, bad, pain, much	1. *"Cuts, blood and then scars, it just makes the pain go away even for a little while. It tells me I'm still alive and can still feel pain"* 2. *"I want to cut my self I use to be like this I wanna feel pain on my wrist."*
Topic 3: Appreciative users	get, really, video, well, use, understand, thing, could, back, go	1. *"I really appreciate that you made this video I am always so worried if I can show my scars or not but this helped a lot so thanks"* 2. *"This video literally changed my mentality so much on what I think off my scars wow, thank you I really needed to hear this from a different perspective."*

As seen in Table 2, the first column represents the topics that describe the entire set of comments: (1) *self-harming users,* (2) *clean (recovered) users*, (3) *at-risk audiences* and (4) *appreciative users.* In interpreting these topics, we grouped viewers based on similar themes. In doing so, this study identified the most discussed topics, as shown in Figure 7. The clean (recovered) comments topic represents the most dominant topic of the discussion. This result is similar to previous research investigating responses to DSH videos on YouTube [12]. Topics about self-harm experiences and at-risk users were discussed at the same rate, while a few comments comprised a topic indicating appreciative responses.

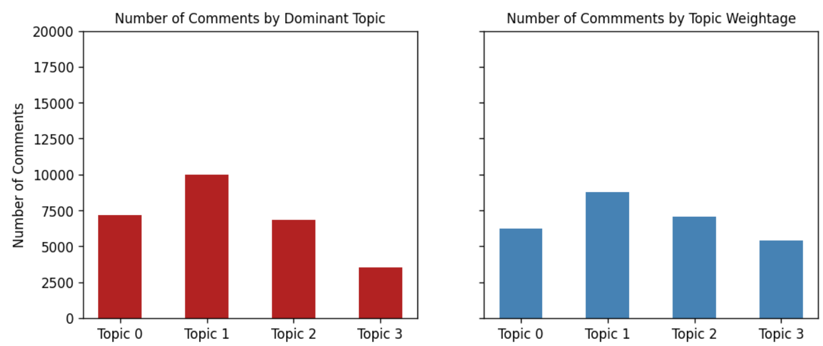

Fig. 7. Dominant topic of the discussion

4.3 Viewers' opinions

Unlike other social media platforms such as Twitter, YouTube does not limit the length of viewers' comments. This allows broad audiences to discuss several topics about the video and express their opinions. However, the present study applied VADER to compute the sentiment score of each comment. In other words, our experiment examined the rate of sentiments from the various topics discovered from the video responses. Figure 8 illustrates the percentage of positive, negative and neutral sentiments of each topic. As seen from the horizontal (colour-coded) bar graph, green represents a set of positive comments, while the yellow colour denotes neutral responses. Additionally, the red colour indicates the family of negative comments.

Notwithstanding, our investigation discovered mixed opinions across the topics. In the first topic, most of the viewers expressed negative opinions as opposed to positive sentiments. This could be a result of their discussion concerning self-harm experiences and associated difficult feelings. This is similar to the third topic, which consisted of more negative sentiments. While only a small number of viewers responded positively, some audiences remained neutral on both topics. In the same vein, topics two and four displayed a significant number of positive opinions. This shows how viewers in our sample encouraged positive peer support by inspiring others with their own recovery process. Also, they conveyed a positive tone in thanking the uploaders and acknowledging that the video contents were informative.

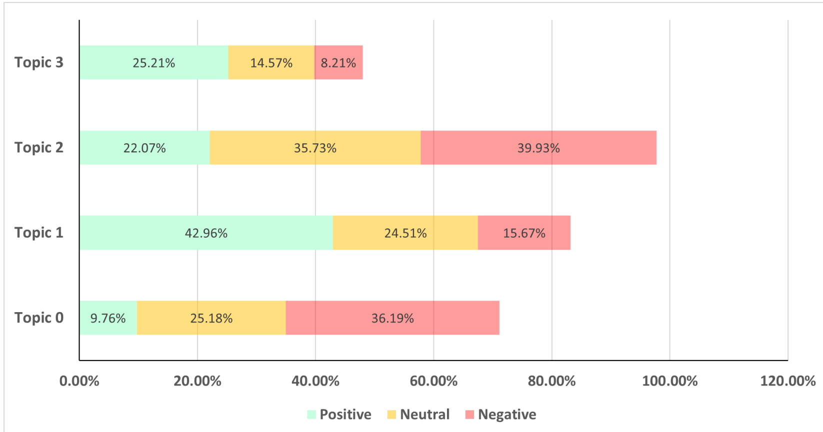

Fig. 8. Topics' sentiment analysis

5 Discussion

The objective of the present study was to investigate the group of people involved in sharing videos about self-harm on YouTube, viewers' opinions on the videos, and the nature of the discussions concerning those videos. Our findings demonstrated that only a small proportion of medical experts and professional bodies reach out to people struggling with self-harm on YouTube. Additionally, videos uploaded by these groups received fewer interactions among viewers and were less favourably viewed by the audiences. This indicates that healthcare and academic professionals need to reconsider the nature of video content preferred by self-injurers, as they favour entertaining videos over videos for educational purposes. The present investigation shows that individuals expressed their self-harming methods or behaviours, including the part of their body they injured, while commenting on the videos in our sample.

Using a statistical topic modelling technique, this study uncovered different topics from 27,520 comments on videos presenting self-harm information. Our analysis grouped viewers who participated in the conversation based on similar topics as one of the following: (1) *self-harming users,* (2) *recovered (clean) viewers*, (3) *at-risk audiences*, and (4) *appreciative users*. Moreover, based on the analysed video comments, many of those who participated in self-harm discussions engaged in sharing ideas and ways to recover from DSH. Existing research has demonstrated that self-harming individuals disclose their personal information on the internet to validate and accept their experiences [14] and receive peer support from members of the online community who are self-harming [13].

On the other hand, we found comments indicating some people accept DSH and normalise the behaviour while putting themselves at risk of potential harm. This finding is supported by existing studies that have explored self-harm content online [13]. Additionally, the comments from *appreciative users*, in which viewers were thanking the YouTuber for sharing a video, shows viewers found helpful information in the video. This is further supported by the high positive sentiments discovered in this group. Therefore, when *at-risk users* received helpful responses addressing dangerous behaviour, the information contained in the comment could be significant and therefore lead them toward recovery, especially if it is from a peer with relevant experience.

Furthermore, viewers had different opinions across the topics. Responses from recovered and appreciative themes attract more positive sentiments in contrast to comments from the topics representing at-risk and self-harming audiences. This result is similar to that of a study which investigated self-harm content on Flickr [28]. The sentiment cues from these topics and other related features can be used to build a machine learning classifier to detect users at risk of potential harm. Hence, to successfully target children and young self-harming individuals who use social media like YouTube and view videos about self-harm, it is essential to consider the content of such videos and their *positive* impact in an online community. One example could be providing innovative digital features to support vulnerable viewers and encourage help-seeking and recovery.

Although social media reveals the voices of people who self-harm, it is equally essential to consider self-harming users when designing and building digital interventions. Consequently, our findings add to the existing studies which demonstrate how young people who engage in DSH prefer YouTube as one of the online social spaces where they access information about their behaviour, interact with their peers with related issues, and exchange support [13].

6 Limitations

The current study has certain limitations. Firstly, our study retrieved and examined 107 videos discussing self-harm on YouTube. These videos do not represent the entire set of videos disseminating self-harm information on YouTube. Although the codebook used in categorising uploaders was reliable, it could be argued that the codebook may not be generalised, as it depicts the view of the researchers' practice. It is also likely that the categories we uncovered do not represent all YouTubers posting videos about self-harm; more research is needed to investigate other hidden categories or sub-categories.

Even though there are several ways to assess the impact of videos about self-harm, investigating the responses to those videos could be an effective strategy to determine the impact of those videos. While our analysis focused on only comments from the examined videos, it is apparent that not all viewers commented on the videos, based on the number of views compared to the number of comments. Therefore, the opinions of individuals who viewed without commenting could be another direction to explore. Secondly, the data we retrieved from YouTube does not incorporate demographic information of viewers who commented on the examined videos. Therefore, it is unknown whether most of those who comments are children, youth, or adults. However, we assume that most were young because of the existing evidence on youth internet use [26].

7 Conclusions and Future Work

Social media spaces like YouTube play a crucial role in disseminating information about mental health and well-being. The present research adds to the body of knowledge regarding the impact of social media concerning self-harm information. While other studies in this field concluded that social media could normalise or reinforce self-harm [15], our study shows the positive influence of social media content in increasing awareness about self-harm and promoting peer support. In response to our first research question, this study examined several channels and sources of YouTube videos presenting information about self-harm. Most of these videos were uploaded (and assigned to the *people and blogs* category) by the group of *non-professionals*. While audiences favourably responded to videos from experts and government institutions, there is a need for stakeholders to increase participation and engagement on social media to support vulnerable users at risk of DSH and consider how to design integrated online interventions.

We addressed the second and third research questions using natural language processing and computational techniques, specifically LDA and VADER. We examined topics from 27,520 responses of 107 videos presenting information about self-harm on YouTube. While viewers who participated in the discussion were grouped based on similar topics, our study found the following themes that described the entire comment corpus: (1)*self-harming users*, (2) *clean (recovered) viewers*, (3) *at-risk audiences* and (4) *appreciative users*. Additionally, topics representing *clean (recovered) viewers* was the most dominant theme among the discussions. The third research question was addressed by analysing the sentiment score of comments on each topic. Sentiment analysis revealed that *clean (recovered)* viewers and appreciative comments contained more positive sentiments than *self-harm* and *at-risk* comments. This is similar to findings of a previous study that examined Flickr content [28].

In addition to looking at other features, future studies should consider incorporating these sentiment-based features into a machine learning classifier to detect self-injurious or at-risk comments. Additionally, there is a need to determine the most effective online methods for reaching adolescents who self-injure, as young people prefer to acquire health information online and, when requested, prefer to receive NSSI assistance online. Medical clinicians supporting self-harming people need to consider investigating how self-harmers source or access information on social media during therapy sessions in order to help redirect their online activity toward positive content and interactions.

References

1. Arendt, F., Scherr, S., Romer, D.: Effects of exposure to self-harm on social media: Evidence from a two-wave panel study among young adults. New Media & Society **21**(11-12), 2422–2442 (2019)
2. Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent dirichlet allocation. the Journal of machine Learning research **3**, 993–1022 (2003)

3. Brunner, R., Kaess, M., Parzer, P., Fischer, G., Carli, V., Hoven, C.W., Wasserman, C., Sarchiapone, M., Resch, F., Apter, A., et al.: Life-time prevalence and psychosocial correlates of adolescent direct self-injurious behavior: A comparative study of findings in 11 european countries. Journal of Child Psychology and Psychiatry **55**(4), 337–348 (2014)
4. Cash, S.J., Thelwall, M., Peck, S.N., Ferrell, J.Z., Bridge, J.A.: Adolescent suicide statements on myspace. Cyberpsychology, Behavior, and Social Networking **16**(3), 166–174 (2013)
5. Cosley, D., Forte, A., Ciolfi, L., McDonald, D.: Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing. ACM Press (2015)
6. De Choudhury, M., Counts, S., Horvitz, E.: Social media as a measurement tool of depression in populations. In: Proceedings of the 5th annual ACM web science conference. pp. 47–56 (2013)
7. De Choudhury, M., Gamon, M., Counts, S., Horvitz, E.: Predicting depression via social media. In: Seventh international AAAI conference on weblogs and social media (2013)
8. Emma Hilton, C.: Unveiling self-harm behaviour: what can social media site twitter tell us about self-harm? a qualitative exploration. Journal of clinical nursing **26**(11-12), 1690–1704 (2017)
9. Hutto, C., Gilbert, E.: Vader: A parsimonious rule-based model for sentiment analysis of social media text. In: Proceedings of the International AAAI Conference on Web and Social Media. vol. 8 (2014)
10. Jacobi, C., Van Atteveldt, W., Welbers, K.: Quantitative analysis of large amounts of journalistic texts using topic modelling. Digital journalism **4**(1), 89–106 (2016)
11. Lehr, C.A., Tan, C.S., Ysseldyke, J.: Alternative schools: A synthesis of state-level policy and research. Remedial and Special Education **30**(1), 19–32 (2009)
12. Lewis, S.P., Heath, N.L., Sornberger, M.J., Arbuthnott, A.E.: Helpful or harmful? an examination of viewers' responses to nonsuicidal self-injury videos on youtube. Journal of Adolescent Health **51**(4), 380–385 (2012)
13. Lewis, S.P., Heath, N.L., St Denis, J.M., Noble, R.: The scope of nonsuicidal self-injury on youtube. Pediatrics **127**(3), e552–e557 (2011)
14. Lewis, S.P., Rosenrot, S.A., Messner, M.A.: Seeking validation in unlikely places: the nature of online questions about non-suicidal self-injury. Archives of Suicide Research **16**(3), 263–272 (2012)
15. Lewis, S.P., Seko, Y.: A double-edged sword: A review of benefits and risks of online nonsuicidal self-injury activities. Journal of clinical psychology **72**(3), 249–262 (2016)
16. Madge, N., Hewitt, A., Hawton, K., Wilde, E.J.d., Corcoran, P., Fekete, S., Heeringen, K.v., Leo, D.D., Ystgaard, M.: Deliberate self-harm within an international community sample of young people: comparative findings from the child & adolescent self-harm in europe (case) study. Journal of child Psychology and Psychiatry **49**(6), 667–677 (2008)
17. Marchant, A., Hawton, K., Stewart, A., Montgomery, P., Singaravelu, V., Lloyd, K., Purdy, N., Daine, K., John, A.: A systematic review of the relationship between internet use, self-harm and suicidal behaviour in young people: The good, the bad and the unknown. PLoS one **12**(8), e0181722 (2017)
18. McDermott, E., Roen, K.: Youth on the virtual edge: Researching marginalized sexualities and genders online. Qualitative Health Research **22**(4), 560–570 (2012)

19. McRoberts, S., Bonsignore, E., Peyton, T., Yarosh, S.: Do it for the viewers! audience engagement behaviors of young youtubers. In: Proceedings of the The 15th International Conference on Interaction Design and Children. pp. 334–343 (2016)
20. Mergel, I.: A framework for interpreting social media interactions in the public sector. Government information quarterly **30**(4), 327–334 (2013)
21. Miguel, E.M., Chou, T., Golik, A., Cornacchio, D., Sanchez, A.L., DeSerisy, M., Comer, J.S.: Examining the scope and patterns of deliberate self-injurious cutting content in popular social media. Depression and anxiety **34**(9), 786–793 (2017)
22. Mitchell, K.J., Ybarra, M.L.: Online behavior of youth who engage in self-harm provides clues for preventive intervention. Preventive medicine **45**(5), 392–396 (2007)
23. Nock, M.K.: Self-injury. Annual review of clinical psychology **6**, 339–363 (2010)
24. Nock, M.K., Holmberg, E.B., Photos, V.I., Michel, B.D.: Self-injurious thoughts and behaviors interview: development, reliability, and validity in an adolescent sample. (2007)
25. Pater, J.A., Haimson, O.L., Andalibi, N., Mynatt, E.D.: “hunger hurts but starving works” characterizing the presentation of eating disorders online. In: Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing. pp. 1185–1200 (2016)
26. Rideout, V.J., Foehr, U.G., Roberts, D.F.: Generation m 2: Media in the lives of 8-to 18-year-olds. Henry J. Kaiser Family Foundation (2010)
27. Rodriquez, J.: Narrating dementia: Self and community in an online forum. Qualitative Health Research **23**(9), 1215–1227 (2013)
28. Wang, Y., Tang, J., Li, J., Li, B., Wan, Y., Mellina, C., O’Hare, N., Chang, Y.: Understanding and discovering deliberate self-harm content in social media. In: Proceedings of the 26th International Conference on World Wide Web. pp. 93–102 (2017)