International Journal of Communication 15(2021), 1244–1262 1932–8036/20210005

Fighting Zika With Honey: An Analysis of YouTube's Video Recommendations on Brazilian YouTube

JONAS KAISER[1]
Harvard University, USA
Alexander von Humboldt Institute for Internet and Society, Germany
Suffolk University, USA

ADRIAN RAUCHFLEISCH[2]
National Taiwan University, Taiwan

YASODARA CÓRDOVA
Harvard University, USA

This study focuses on misinformation about the Zika virus on YouTube in Brazil. Whereas most studies in health communication have so far conducted manual content analysis to understand the scope of health-related misinformation, we use a computational approach to understand the formation of information clusters via YouTube's video recommendation algorithm on a platform level. Through network analysis and topic modeling of 20,499 YouTube videos, we are able to show a clear distinction between Portuguese-language and English-language YouTube recommendations in regard to this topic. The most prominent videos of both communities feature mostly accurate and correct information about Zika, but both communities are prone to faulty and potentially dangerous misinformation. We show that this harmful misinformation is not separate from the communities that feature correct information. In addition, we suggest that health misinformation surrounding both Zika and vaccinations seems to be more prevalent in YouTube's video recommendation long tail.

Keywords: YouTube, algorithm, Brazil, health communication, misinformation

The Zika virus first was identified in Brazil in mid-2015. By the end of 2016, Brazil had reported more than 205,000 Zika cases (Lowe et al., 2018). Because the virus is transmitted by mosquitoes, blood

---

Jonas Kaiser: jkaiser@cyber.harvard.edu
Adrian Rauchfleisch: adrian.rauchfleisch@gmail.com
Yasodara Córdova: who@yaso.is
Date submitted: 2020–06–15

[1] The authors want to thank the anonymous reviewers for their helpful comments and suggestions.
[2] Adrian Rauchfleisch's research was funded by the Ministry of Science and Technology, Taiwan (Grant No 108-2410-H-002–007-MY2).

transfusions, and sexual contact, and it can infect babies still in the womb and cause microcephaly, having reliable information is paramount. Although Brazilians know the most important facts about the Zika virus, misinformation and conspiracy theories are still widespread in Brazil (Carey, Chi, Flynn, Nyhan, & Zeitzoff, 2020). One of the vectors for the widespread belief could be YouTube, as Bora, Das, Barman, and Borah (2018) show in their study about Zika videos on YouTube; almost one quarter of the analyzed videos about Zika in their study were classified as misleading. Our study connects with this research as we analyze the role of health misinformation around Zika on Brazilian YouTube, with a main focus on the video recommendations on the platform.

There is a trove of research on misinformation, disinformation, or "fake news" in the political—often Western European and U.S. American—context (Tucker et al., 2018). However, there is less research on the role that misinformation plays in other fields that are primarily nonpolitical, such as health communication—especially in countries in the Global South, such as Brazil. Furthermore, most research on misinformation focuses either on social media platforms, such as Twitter or Facebook (e.g., Guess, Nagler, & Tucker, 2019), or on exposure and/or sharing of news (e.g., Guess, Nyhan & Reifler, 2018). YouTube, in comparison with other platforms, has been mostly overlooked research-wise, although scholars are slowly catching up. Recently, the platform's recommendation algorithms have been put under special scrutiny (e.g., Kaiser & Rauchfleisch, forthcoming; Kaiser & Rauchfleisch, 2018). In a report by *The New York Times*, for example, YouTube stated publicly that about 70% of video views came from its video recommendations (Fisher & Taub, 2019).

Against this background, this case study's focus is the Zika virus, which is primarily spread through mosquito bites and caused a public health emergency in Brazil in 2016; the virus was still active in 2019 (Jacobs, 2019). Because 42% of Brazilians (Newman, Fletcher, Kalogeropoulos, & Nielsen, 2019) get their news from the Alphabet-owned video platform, it is imperative to understand what type of content a simple search for "Zika" will display on YouTube and how the platform's recommendation algorithms contribute to the spread of health misinformation. In our study, we specifically focus on the recommendations seen by users when they just check YouTube for Zika-related information without being logged in to their YouTube account. Although this is a limitation of our study, it allows us to identify generalized "platform-level recommendations" (Roth, Mazières, & Menezes, 2020) that most likely apply globally to all users for the existing videos about Zika on YouTube. We thus extend the rather small-scale analysis from Bora et al. (2018) with a large-scale analysis of the video recommendations on the platform. In our analysis of YouTube's video recommendation algorithm, we show that there is a difference between the most prominent videos on Zika and the long tail of recommended videos: Whereas the former consists of mostly legitimate channels and information about Zika, the latter contains many videos from channels that promote conspiracy theories or alternative healing methods. Finally, we discuss explanations and solutions.

Misinformation in Brazil

In the aftermath of the 2016 U.S. presidential election, the academic study of the spread of misinformation has gained prominence, eventually even forming its own research field (e.g., Benkler, Faris, & Roberts, 2018; Freelon & Wells, 2020; Guo & Vargo, 2020; Kaiser, Rauchfleisch, & Bourassa, 2020; Lewandowsky, Ecker, & Cook, 2017). Although initially, the umbrella term for incorrect information was *fake* *news*, academics quickly started differentiating between different kinds of false information. Wardle (2018) was one of the first to create a typology, distinguishing between satire or parody, misleading content, imposter content, fabricated content, false connections, false content, and manipulated content. In her typology, Wardle firmly focuses on the content itself and the intention behind the content creation. With her approach, she is able to separate between different types of fake news and, in a next step, differentiates between misinformation, disinformation, and malinformation. But, as other scholars have noted, there are more forms of creating, distributing, and sharing falsehoods. As Egelhofer and Lecheler (2019) note in their literature review, a broader distinction would be among misinformation, disinformation, rumors, conspiracies, and propaganda. This broader typology is similar to Wardle's, based on the communicator's intent, or, as Freelon and Wells (2020) call it, "the cognitive domain" (p. 5). The difference between disinformation and misinformation, then, is that disinformation describes the deliberate effort to create or share false information with the goal of deceiving the audience, whereas misinformation is the unknowing creation or sharing of false information. Fake news, in this sense, would be a subcategory of disinformation (Egelhofer & Lecheler, 2019).

Freelon and Wells (2020) add to this literature by highlighting that there are two distinct strands of research on disinformation: one looking at the content, and one looking at the reception. The former asks for the content of mis- or disinformation and how it might have been disseminated, and the latter asks for the audience's response to it. Although the field is still in its infancy (Freelon & Wells, 2020), researchers have made progress in shedding light on the issue. For instance, Bradshaw et al. (2020) can show that in the case of the U.S., people on Twitter shared "junk news" as often as regular news. Similarly, as Chaves and Brage (2019) show, Brazilian fact-checking sites had to post "228 verifications of false stories disseminated on social media and/or messaging apps, covering a range of about 132 different topics" (p. 474) during the 20 days between the first and second rounds of the 2018 presidential election. Finally, a study on the effect of fake news in Germany by Zimmermann and Kohring (2020) finds that people who trust neither the political system nor the media are more prone to believing disinformation. This is especially relevant given that a 2018 Gallup poll showed that only 17% of Brazilians were confident in their government (Reinhart, 2018). In addition, the Reuters Digital News Report (Newman et al., 2019) found that Brazilians' trust in the media dropped from 59% in 2018 to 48% in 2019—a drop that the researchers attribute to the "political atmosphere." In comparison, 31% of Brazilians trust news from social media such as YouTube.

In this article, we follow Egelhofer and Lecheler's (2019) definition and thus will only use the term *misinformation* going forward; we do not want to speculate about a video creator's intent. This is based on two aspects: Because we are primarily interested in the communities that YouTube's algorithms form and thus not in conducting a content analysis (cf. Bora et al., 2018), we can speak about only the content of some specific examples, not the whole sample. Consequently, we cannot draw conclusions about the video creators' intent and thus prefer to assume good faith. Indeed, it is possible that the video creators who tout garlic or honey as a cure for Zika simply do not know better. Yet, our analysis falls in the research area that Freelon and Wells (2020) call content as we look at the overall topics of YouTube videos. However, we do not conduct a manual content analysis to see whether a video contains misinformation; rather, we analyze the topical communities that YouTube's video algorithm forms to see whether distinct communities of misinformation are being created. We thus aim to contribute to the literature on misinformation in two ways. First, we will show how YouTube's algorithms can potentially contribute to the spread of misinformation.

Second, we will add to the literature by focusing on Brazil, a country that is both understudied in that regard and potentially less closely moderated than the U.S. or Western European countries.

YouTube

More than 2 billion unique users visit YouTube on a monthly basis (YouTube, 2020); consequently, it is an important cornerstone in the networked public sphere (Benkler, 2006) because it gives its users the ability not only to circumvent the traditional gatekeepers, but also to share its content on other social media platforms or websites. According to Kim (2012), YouTube can best be described as “a convergence medium between the Internet and TV” that highlights “a series of contradictions between traditional broadcasting and digital narrowcasting” (p. 53). YouTube is part of an “expansive ecosystem of connective media” (van Dijck & Poell, 2013, p. 5) in which users have the ability to produce their own content (Burgess & Greenberg, 2009). Although this feature of the platform leads to positive outcomes, such as critical citizen journalism (Antony & Thomas, 2010), it can also lead to the spreading of rather questionable content, such as antivaccination messages (Briones, Nan, Madden, & Waks, 2012) or conspiracy theories about climate change (Allgaier, 2019). In our research, we were able to show that communities on YouTube can form around conspiracy theories (Rauchfleisch & Kaiser, 2020).

As stated earlier, 42% of Brazilians are using YouTube for news (Newman et al., 2019). However, the video platform is used not only for news, but also for all sorts of entertainment (e.g., gaming, music, sports) and information (e.g., how to build computers or to do makeup), including health information, such as on Zika or COVID-19. In general, YouTube has become an important platform for science- and health-related information, especially for young people (for an overview, see Allgaier, 2020). In Germany, for example, 42% of respondents between 14 and 29 years of age stated in a national survey that they regularly search science-related information on YouTube (Wissenschaft im Dialog, 2018), highlighting both the integral part the platform plays in people’s lives and the need for researchers to shed light on what is happening on the platform and how YouTube’s algorithms contribute to these activities.

YouTube and Health Misinformation

Many studies in different disciplinary contexts have already focused on misinformation. Especially in the medical and health context, researchers have analyzed the content on video platforms such as YouTube. One strand of research in this area usually analyzes all relevant videos identified with specific search terms and then perform a manual content analysis of the videos (for an overview, see Gabarron, Fernandez-Luque, Armayones, & Lau, 2013). Syed-Abdul et al. (2013), for example, analyzed videos about anorexia on YouTube and came to the conclusion that misleading pro-anorexia videos were overall the most popular videos. In a similar analysis, Pant et al. (2012) evaluated YouTube videos about acute myocardial infarction. They concluded that only a few videos from reputable sources were available and that these videos were usually not very popular. However, it seems to depend on the context. In their analysis, Sood, Sarangi, Pandey, and Murugiah (2011) classified as helpful more than half of the available YouTube videos about kidney stones. Besides medical issues, studies have also focused on science-related issues that are often targeted by conspiracy theories. In his analysis of YouTube videos about climate change, Allgaier (2019) found that almost half of the videos in his analysis promote conspiracy theories about climate change denial. To our knowledge, no studies have been conducted on the role that the algorithms play in potentially promoting misleading or harmful health information.

Because of the Zika outbreak in 2016, some researchers also studied the information and misinformation that users will find on YouTube. Indeed, in an analysis on the most popular videos on Zika, Nerghes, Kerkhof, and Hellsten (2018) showed that about one third of the videos contained misinformation, and there was no significant difference in how users interacted with legitimate videos as compared with those that contained misinformation. In a similar analysis, Bora et al. (2018) found that although there were more legitimate videos (~70%) on Zika in the top videos, the videos containing misinformation had more views, likes, and shares. We intend to add to this literature by focusing specifically on Brazil and on the networks that the recommendation algorithm creates.

Research Questions

With regard to the platform architecture, most studies have only focused on the content of a specific subset of videos. Other elements, such as the video recommendations, were not the main focus so far of studies in the area of health and science communication. Therefore, we focus on YouTube's video recommendations on Brazilian YouTube in the case of Zika. In doing so, we deliberately approach the issue differently than other scholars (e.g., Bora et al., 2018) who have conducted a content analysis of the top search results. In prioritizing the recommendation algorithms, we are less interested in specific cases of misinformation and more interested in whether the algorithm creates, for example, specific video communities on the platform level that are united by spreading falsehoods (e.g., antivaccination content) and that can be identified through computational methods on a large scale. If we were to find a community that mostly consisted of videos promoting misinformation, this would suggest that YouTube's algorithms could differentiate between true and false; thus, YouTube would be able to remove the misleading videos easily. A study by O'Callaghan, Greene, Conway, Carthy, and Cunningham (2015), for example, was able to establish that YouTube had created a far-right extremist filter bubble. If, however, we found that these video communities did not exist, and misinforming videos were in the same communities as factual videos (i.e., were being recommended in the same context as factual videos), this would paint a much more complicated picture in which true information and false information are only clicks apart on YouTube. In this sense, we are specifically interested in this question:

RQ1: What video network communities does YouTube's video recommendation algorithm form in the case of Zika?

Given that video recommendation networks often form "noisy" communities that are due to YouTube recommending a variety of nevertheless related content, we are especially interested in the underlying patterns that drive this "noise" and how these inform the video network communities that we identify in RQ1. A way to identify and understand these patterns is to know what topics the videos deal with. Because Zika has been associated with conspiracy theories, such as one suggesting that Zika is caused by vaccines or that Zika was created by "genetically modified mosquitoes" (Klofstad, Uscinski, Connolly, & West, 2019), understanding the specific topics might help uncover algorithmic communities of misinformation. We thus ask:

RQ2: What topic communities can be found on YouTube with regard to Zika?

Finally, we are interested in the stage at which videos that might contain misinformation get recommended. In its quest to keep its own platform safe, YouTube might, for example, actively curate the initial search results that users will see when searching for "Zika"; this raises the question of whether or not the video recommendations for these search results are also sanitized:

RQ3: What are the differences in the flow of recommendations between different topics?

In short, this question aims at understanding differences in topic recommendations and, in general, whether clicking on a video and then another will keep the user on topic (i.e., videos about Zika) or expose the user to different topics.

Method

For our analysis, we collected the top 450 results of a search for "Zika" on Brazilian YouTube. We created a scraper on a Brazilian server for our search, and we set Portuguese as the interface language and Brazil as the location in a nonpersonalized headless browser. We then collected the top 10 videos that YouTube recommended alongside these videos with YouTube's API. We chose the top 10 to simulate the behavior of users who do not have a YouTube account and will thus receive less personalized results. Although we could have collected even more recommended videos, we found it unlikely that users would scroll to the bottom of a video's page to see all recommendations.[3] We then repeated this step once more for the videos that were added. Although the lack of personalization for the recommendations (cf. Covington, Adams, & Sargin, 2016; Zhao et al., 2019) could be a potential limitation—which, of course, we acknowledge—it is less of an issue in our case because we were mainly interested in what users see when they only check on YouTube for Zika-related information without being logged in to their YouTube account. The personalization has an impact on the video recommendations as soon as a user is logged in and has the personalization turned on. Our analysis provides thus a general baseline and an overview of the existing videos about Zika on YouTube. Roth et al. (2020) propose a similar approach with a focus on so-called platform-level recommendations that most likely apply globally to all users. This helps us to better understand the navigation topology and, more specifically, the video recommendation algorithm that is a central element of YouTube's platform architecture.

To be able to understand not only how videos were related to each other through YouTube's recommendation algorithm, but also the underlying structure, we aggregated the videos to their respective channels. In doing so, we have networks for both videos and channels where the connecting edges represent YouTube's video recommendation algorithm (see Figure 1). In the channel network, the edges, which represent recommendations from videos of one channel to videos of another, can have weights of more than one in the case of videos of one channel recommending videos of another channel more than once.

---

[3] In its own publicly available studies about the video recommendation algorithm, YouTube does not analyze the role of the rank beyond rank 9 (Zhao et al., 2019).

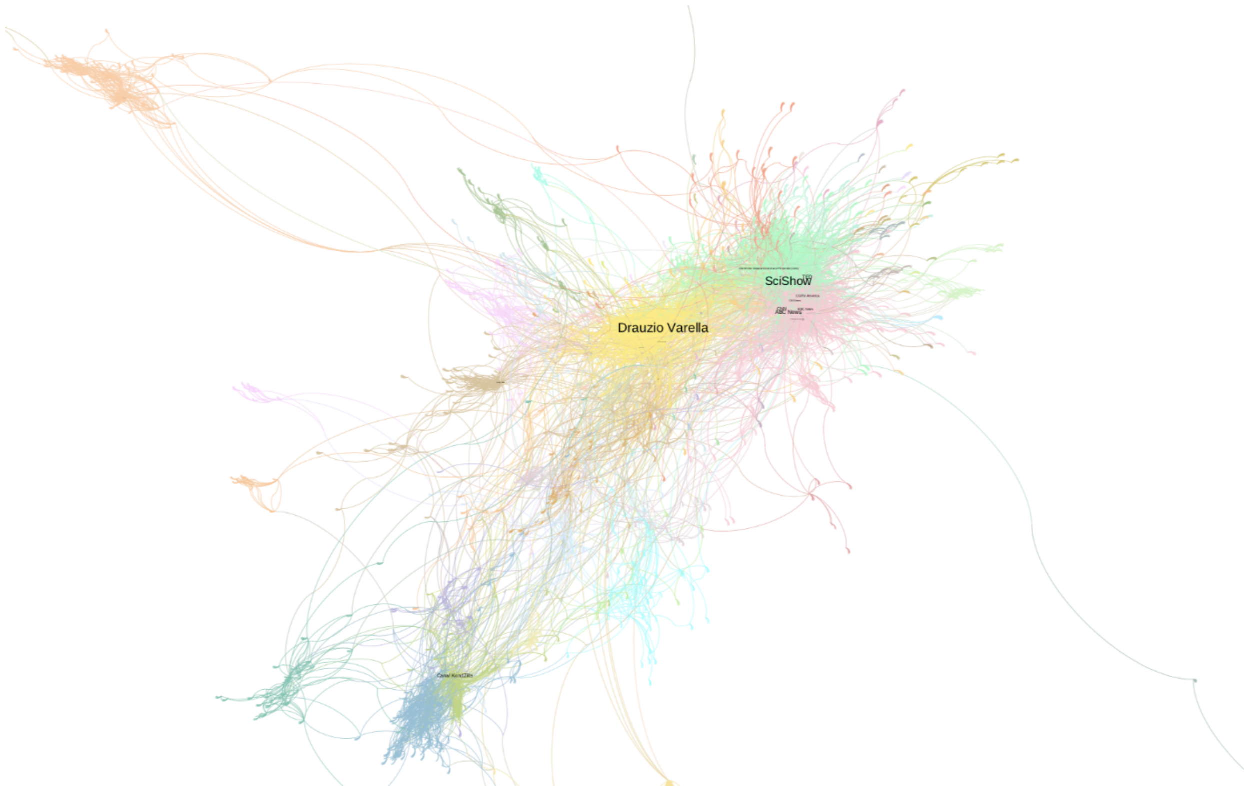

Figure 1. Channel network of Zika video recommendations (nodes = 10,003, edges = 19,522; node size per indegree).

We used the Louvain algorithm implemented by Traag (2015) in Python to identify communities in our directed graphs. We used the Louvain algorithm because of its efficiency, and Traag's (2015) implementation gives stable solutions. In the case of video recommendation graphs, in which nodes represent either videos or channels, the identified communities can be best described as information clusters.[4]

In the next step, we took all video descriptions, including video tags, and video titles and used them as the basis for a topic modeling analysis with the R package "stm" (structural topic model; Roberts, Stewarts, & Tingley, 2019). We used stop-word lists for Portuguese and English, removed all punctuation and numbers, and eventually stemmed the words in English, Spanish, and Portuguese (*n* = 20,499 documents).[5] Furthermore, we only used words that appeared at least 10 times in the whole corpus. We tested the topic with 10, 20, 30, 40, 50, 60, 70, 80, 90, 100, 150, and 200 topics; after checking the words with the highest probability and the most exclusive words (FREX) for each topic, we decided on 200 topics

[4] The modularity score was .91 for the video graph and .82 for the channel graph based on the video recommendations.

[5] Ten videos had to be excluded from the original 20,509 videos because not enough textual information was available to include them in the topic model.

because this gave us the most granular number of relevant health-related topics (n = 12).[6] We also calculated the overlap between videos based on their tags (Figure 2). Tags are chosen by the channel owners and usually are not visible to a user watching a video. We found that 46.8% of the video recommendations (directed edges) had an overlap based on at least one tag (e.g., both videos are tagged with "Zika") for the source and the recommended video; 31.8% of all edges had an overlap on two tags, 23.0% on three tags, 12.4% on at least five tags, and 5.9% on at least 10 tags (Figure 2).

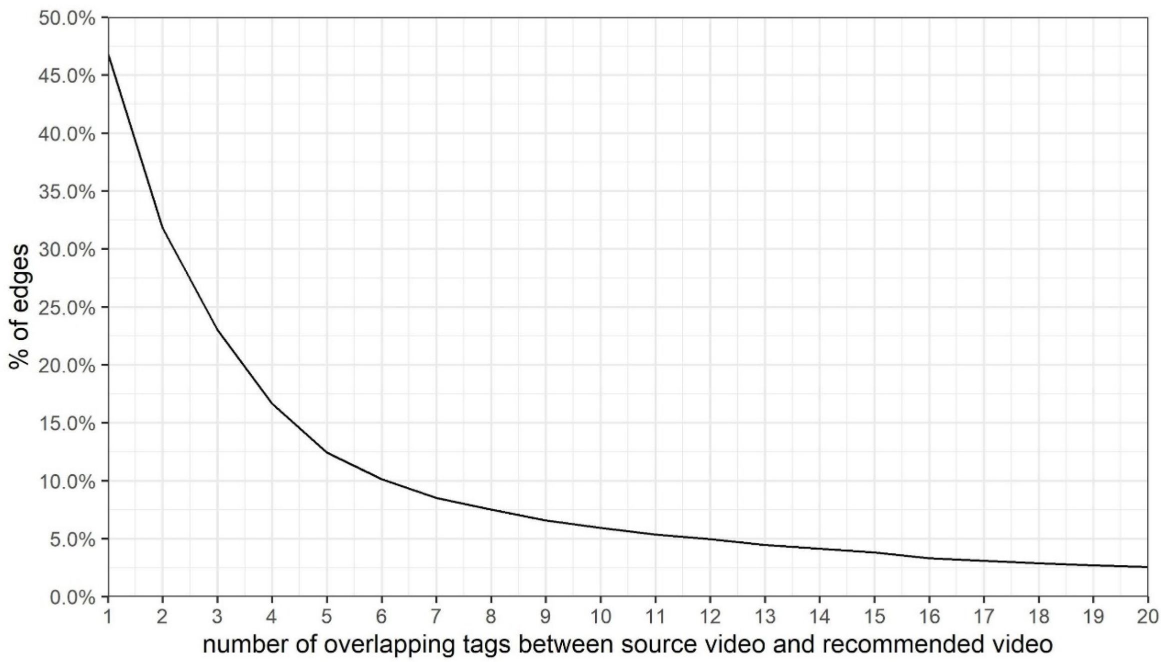

Figure 2. Overlap between videos based on N number of tags.

Results

To make sense of the data we collected, we present our findings in several steps. We first briefly summarize the channel network, which is based on the video recommendations, to give an overview of the channel structure. We then present the video network to answer RQ1. Next, we present the topic model and the associated topic model network to answer RQ2. Finally, we present the Sankey network of topics and the recommendation stages to answer RQ3.

[6] The semantic coherence, the held-out likelihood, the lower bound, and the residuals all indicate that a higher number of topics is better than a number below 150 topics.

Channels

We show that on Brazilian YouTube, the Zika videos can be roughly divided into three communities (Figure 1): Brazilian videos on Zika (the yellow community), English-language videos on Zika (green and pink communities), and videos on the music artist Zika (dark green and blue communities). This by itself highlights YouTube's inherent international character and the diversity of content. However, when analyzing the channel networks (Figure 1), we can differentiate between the most recommended channels (i.e., the channels whose videos got recommended the most) and the so-called long tail, that is, the channels that were also recommended but received fewer recommendations. In our manual analysis of the 20 most recommended videos in our network, we were able to identify only one video that we noted as containing falsehoods. So although the most prominent channels contained mostly correct information on Zika, several of the smaller channels revolved around health misinformation. One video, for example, proposed several home remedies, such as honey, to fight Zika. Indeed, perhaps the biggest finding in this context is that there was no clear delineation between channels that published information on Zika and those that published misinformation on a community detection level.

Videos

When looking at the video level rather than the aggregate channel level, we could see that while the core recommendations of the videos were on the Zika virus, the recommendation algorithms quickly also recommended unrelated content (Figure 3). For example, several music communities in the network were mostly getting recommended through videos from the musician MC Zika and Slovakian YouTuber Lady Zika. Similarly, a notable Serbian and Yugoslavian TV community was in the network. The algorithm's logic, in short, seemed to connect videos on Zika in English to an old Yugoslavian TV show called Žikina dinastija (Žika's Dynasty), and from there to Serbian TV.

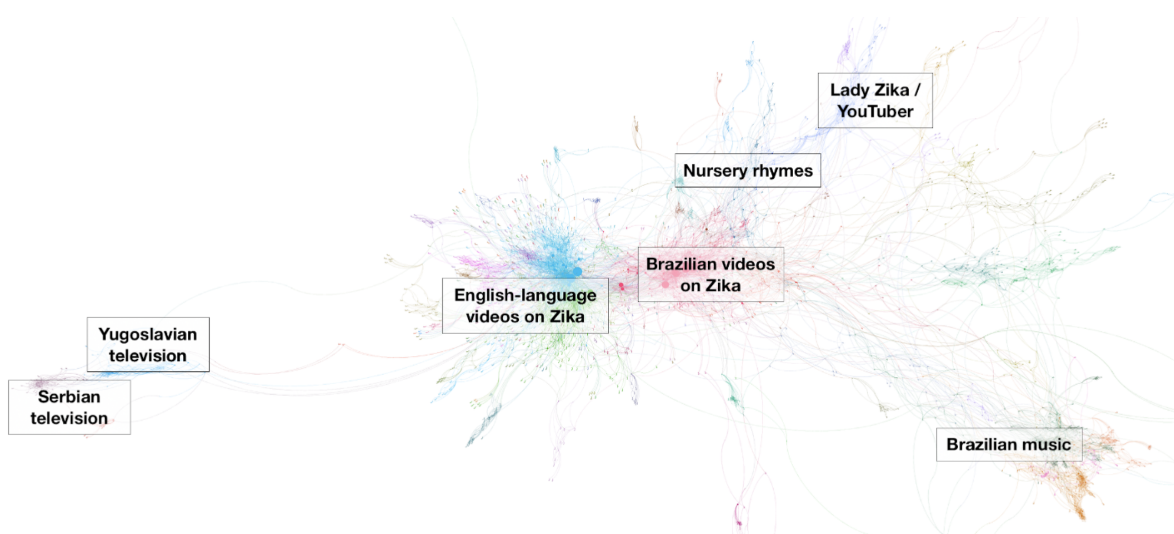

Figure 3. Video recommendation network (nodes = 20,509; edges = 30,583; community detection with Louvain).

When focusing on the video core network—that is, the Brazilian and English-language communities—we saw several English-language communities, but only two core Brazilian communities (Figure 4). A close inspection of the core network based on the video names suggested that the core Brazilian Zika community was divided into one subcommunity that mostly focused on Zika (Figure 4; for example, videos that talked about its origins); and another subcommunity that discussed other mosquito-transmitted diseases, such as dengue or chikungunya (the red community on the right in Figure 4). Some of the most prominent videos in the Brazilian core community, such as *Zika Virus | Origem*, came from Drauzio Varella, a Brazilian doctor, scientist, and successful YouTuber; his channel has 1.97 million subscribers. There was, however, misinformation as well. Several of the less frequently recommended videos in both communities spread misinformation: Some spread myths about where Zika comes from, while others present home remedies for Zika in the form of garlic or honey. The community, then, that connected the Brazilian and English-language communities consisted of videos that talk about Zika in general—such as Zika Virus: What We Know (And What We Don't) from the SciShow channel—but also how mothers and families are dealing with children who have microcephaly, which has been linked to Zika. From this community, YouTube recommended videos on mosquitoes, and more general virus infections such as Ebola. The blue community on the left, then, mostly consisted of Tedx videos (see Figure 4).

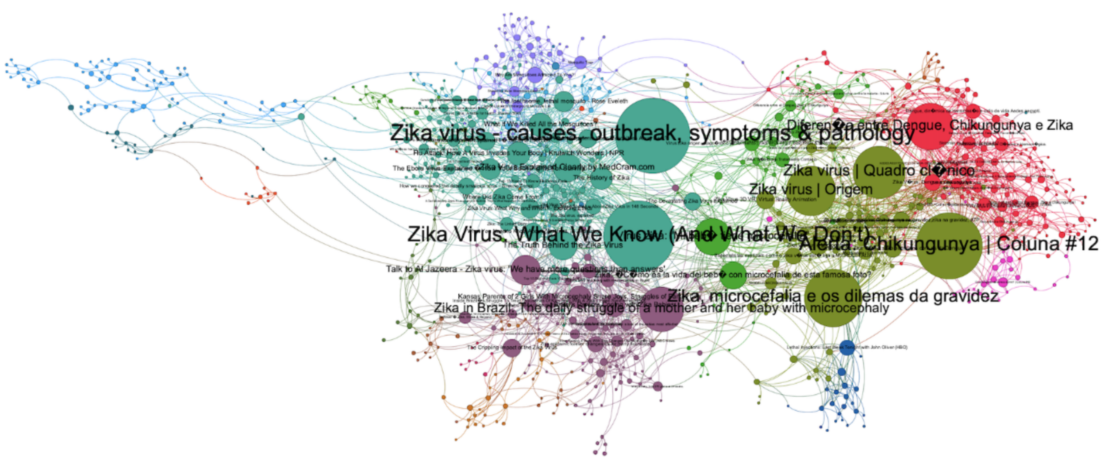

Figure 4. Core of video recommendation network (nodes = 988; edges = 2,888; community detection with Louvain; filter >= 2 indegree, i.e., videos had to be recommended at least twice to appear in the network).

When trying to answer RQ1, we can say that although we were able to identify clear communities in the recommendation network, no community seemed to center on misinformation or conspiracy theories. Instead, we found that the most recommended videos had legitimate creators, such as Drauzio Varella or SciShow. However, this does not mean that misinformation was not present as well; instead, it seems that misinformation on Zika can be found in the same communities in which one also finds legitimate information. This suggests that YouTube's algorithms cannot distinguish between correct and incorrect information.

Topic Modeling

To better understand what topics YouTube's algorithms recommended, we conducted a topic modeling analysis based on all 20,499 video descriptions, video tags, and video titles. We picked 200 topics so that we could identify more granular topics. After carefully looking through all 200 topics, we selected 12 topics that were the most relevant to our case (Table 1). In doing so, we were able to highlight several key aspects: Zika was its own topic in our YouTube network, but we also saw that Zika led to other mosquito-transmitted diseases, such as dengue and chikungunya. At the same time, however, we also found other disease topics, such as cancer (via microcephaly), HIV, and even diabetes. This shows that health videos were closely connected on YouTube through recommendation algorithms, and the topic of mosquitoes was itself likewise grouped. In addition, the search for Zika also presented videos that talked about babies and pregnancy—a finding that makes sense given that Zika is most dangerous for pregnant women. In addition, we found topics on remedies, health, and the immune system, but also (anti-)vaccination viewpoints (some of the videos mock antivaccination talking points, and others have titles such as *New Proof Vaccines for Pharma Profit, Not Health* from RT America).

Table 1. Relevant Topic Labels and Words With the Highest Probability per Topic.

Topic	Words with highest probability (stemmed)
Mosquitoes	mosquit, bug, repel, insect, bit, kill, control
Baby & pregnancy	bab, pregnanc, week, pregnant, mom, updat, trimest
Remedies	saud, natur, agu, cur, vitamin, medic, remedi
(Anti-)Vaccination	vaccin, anti, poli, asthma, flu, autism, health
Health	bod, health, weight, eat, fat, exerc, diet
Immune system	test, cell, type, elis, immun, dna, antibod
Dengue & chikungunya	dengu, mosquit, chikunguny, aed, sintom, aegypt, doenc
Zika	zik, virus, health, mosquit, microcefal, microcephal, brazil
Diabetes	diabet, dor, doenc, ovo, sintom, verm, tratament
Dengue & chikungunya treatment	hom, dengu, fev, rem, treatment, symptom, chikunguny
Cancer	canc, brain, health, tumor, treatment, diabet, cur
HIV	dis, medic, health, hiv, condit, infect, medicin

To understand how the topics that we identified mapped onto the video recommendation network, we identified all videos in the topic model in which one of the 12 chosen topics was the most prominent topic—that is, where it could be assumed that the video was most likely about that topic. We then imported the topics from the topic model into Gephi and colored the nodes based on the different topics (Figure 5). In doing so, we could see that the topics not only matched the identified communities in Figure 4, but also gave more context to what was being discussed in the video core network. More specifically, we could see that, similar to the community detection, the Brazilian videos could be mostly distinguished as those that talked about Zika and associated diseases, and those that presented cures and remedies. The English-language videos, however, were much more fragmented, with fewer videos that led from one topic to another. But, most important, we were not able to identify a misinformation topic here either. As mentioned earlier in the context of the (anti-)vaccination topic, we could find both informing and misinforming videos in this topic community. Indeed, although the long tail of the relevant videos seemed to consist of at least some misinforming videos, these were always connected to the overarching topic communities and, thus, trustworthy videos. Misinformation, in this sense, was always only one click away.

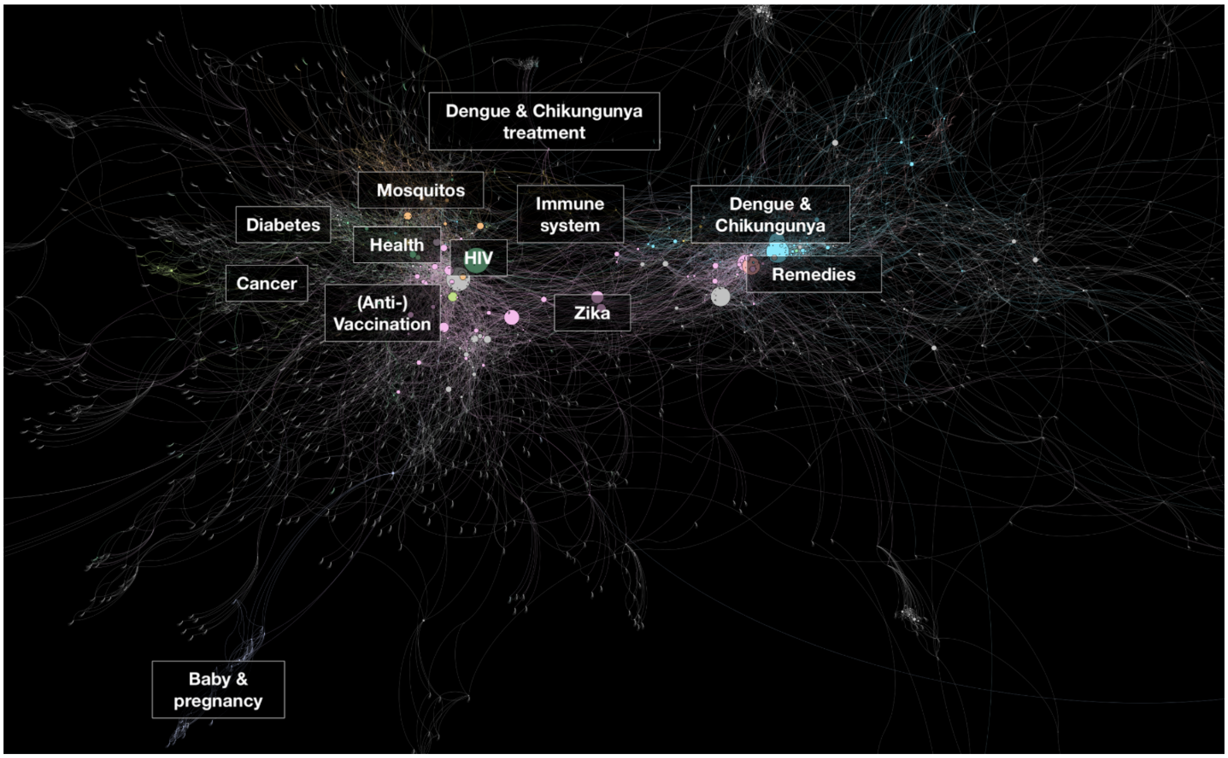

Figure 5. Network of the core video community; node colors and labels represent topics. The 12 relevant topics have been colored according to each topic (11.06%); nodes with no assigned topic were colored gray (88.94%).

Flow of YouTube Recommendations

Finally, we were interested in understanding differences in topic recommendations between (1) search results, (2) the first round video of recommendation, and (3) the second round of video recommendation because this would give us an idea about how YouTube's recommendations have a potential impact on user behavior. To visualize this, we opted for a so-called Sankey network, which visualizes the flow from one variable to another. In our analysis, we can see that the Zika topic became less and less prominent with every recommendation round. Although it was, unsurprisingly, the most prominent topic in the search results, the first and second recommendations saw Zika losing in relative prominence (Figure 6). Indeed, in the first and second recommendation round, the dengue and chikungunya topic moved to the most recommended spot—but even that topic lost relative prominence in the second round. The HIV topic, on the other hand, was barely present in the search results, but became increasingly pronounced with each recommendation step. The same was true for the remedies and baby and pregnancy topics.

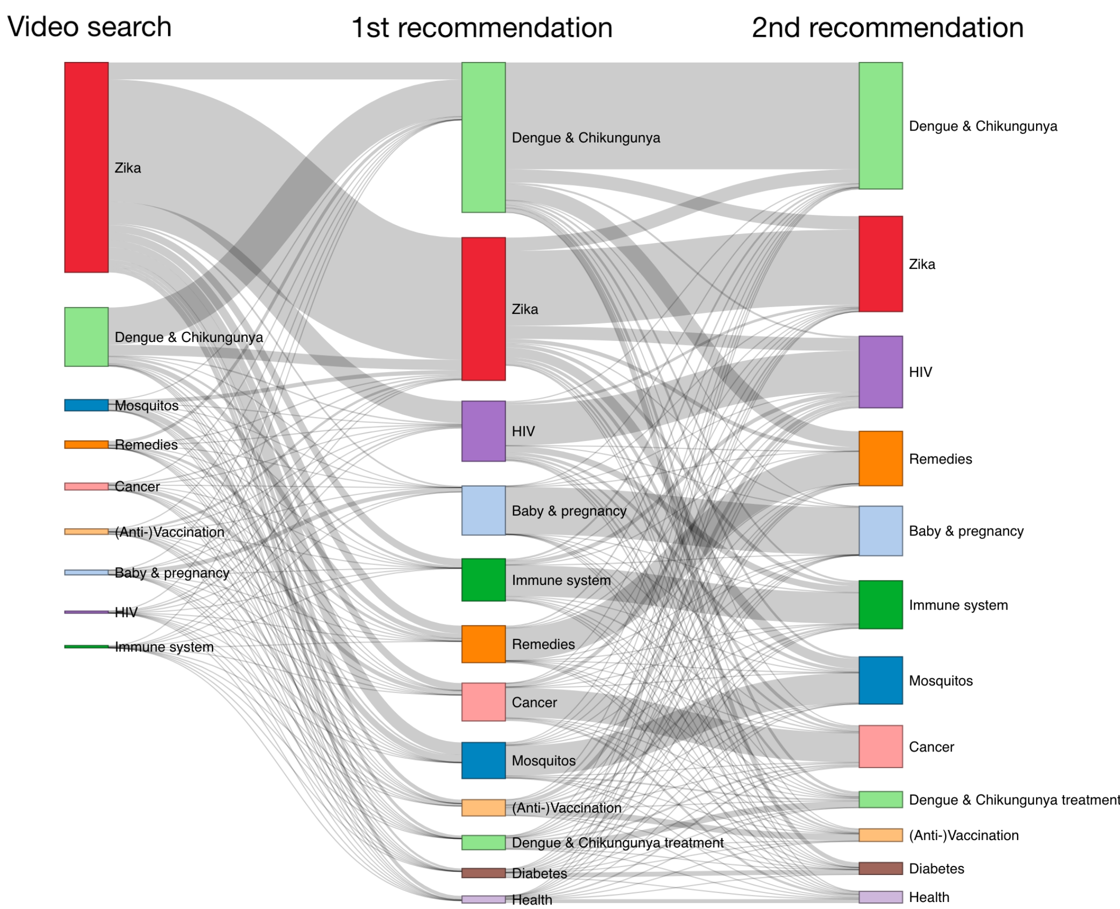

Figure 6. Sankey network of YouTube recommendation flow between topics for three steps (initial video search, first video recommendation, second video recommendation; node color = topics, node size = degree, edges = outdegree; nonrelevant topics were omitted for this visualization).

A more general observation of Figure 6 is that with each recommendation step, the initially prominent topics seem to lose in relative prominence, while initially less prominent or not even existing topics gain in relative prominence. The topics health, diabetes, and dengue and chikungunya treatment, for example, were not present in the initial search results, but then slightly grew from the first round to the second. This could suggest that YouTube's video recommendation system was branching out to other topics in its later recommendation stages—potentially because there were more videos about these topics than more specific issues such as Zika. In addition, it also might mean that with every recommendation, there is a better chance of stumbling on misinformation. The video *Cómo Curar El Virus Del Zika Con Estos 5 Remedios Caseros Milagrosos Y Seguros*, which claims that home remedies such as honey or garlic cure Zika, for example, was added to our network in the second round (YouTube later removed the video). Similarly, the RT America video *New Proof Vaccines for Pharma Profit, Not Health,* which casts doubt on vaccinations, and the video *You'll Be Glad You Watched This Before Vaccinating Your Child!*, which features antivaccination viewpoints (for example, it features the sentence, "In my view, vaccination science is fuzzy science") from a verified YouTube account, were added in the second round of recommendations.

In summary, it seems that recommendations add more topical serendipity and diversity with every step. Topical serendipity in this context refers to the loss of prevalence of the topic that we initially searched for, Zika, and other topics gaining slowly in prominence. Diversity, then, refers to the diversity of viewpoints and content. Although we cannot quantify on which level how much misinformation has been added, it seems that most misinforming videos were added in the first or second round of the recommendations, whereas the initial search results were mostly correct. This suggests that even curating the search results does not protect people from being exposed to misinformation.

Conclusion

In this case study, we showed that health misinformation around Zika is a big issue on YouTube. Although the top video recommendations are mostly trustworthy, not only is there a great deal of misinforming content on YouTube, but it is also often only one click away. Indeed, we show that misinformation around Zika is not isolated to a potential filter bubble full of conspiracy theories; rather, it seems to reside in the recommendation algorithm's long tail. This means that there is a lot of false information on YouTube, and this content will eventually be recommended. At the same time, the identified video communities with the highest number of views focus on music. Although this is partly an artifact of our search term and its intersection with the music artist, it illustrates the typical structure that can be observed on YouTube. Entertainment videos dominate with regard to the number of views, averaging more than a million views per video. Videos that contain misinformation (e.g., RT's antivaccination video = 9,646 views; a video that promotes honey as cure = 120,491 views) that we named in this study fall significantly short in terms of viewership.

That said, it is important to note that this analysis is based on search results for one keyword and its associated video recommendations. Although we tried to avoid personalization at all costs (tested with an API and via browser, public and incognito mode, with and without a VPN, etc.), it is likely that there is some bias in our analysis. We are also intrigued by the finding that there are numerous English-language health-related topics in the model, but only two topics for the Brazilian videos. Although it is possible that the topic model software we used might work better for English, it is unlikely because we have used stm successfully before for languages other than English (Rauchfleisch & Kaiser, 2020). Another more plausible explanation might be that doing our research from a Brazilian server might lead to more fragmented English-language results (potentially because of prominence), which, although connected, are still distinct from each other. This explanation is also supported by the search results. In the search results at the beginning of our analysis, only three videos in the top 20, and 25 videos in the top 100, are in English. However, from rank 101 to 450, English videos were slightly more prominent than Portuguese videos. The lower a video's rank in the search results, the more divergent the content is linguistically from the interface and location settings, as well as topically from the intended interpretation of the search term. This also highlights the double-edged sword that is the recommendation algorithm: Contentious issues often will be targeted by the creators of misinformation, and if there is little content, YouTube will recommend whatever is available, regardless of the source. This is what Golebiewski and boyd (2018) call "data voids." The issue of recommendation algorithms goes beyond that, however. On YouTube, there is no "end" to the recommendations, and thus, bad content will eventually get recommended—data void or not. We saw this in our study: The most prominent search results and recommendations were mostly legitimate; the problem, however, was the long tail. Indeed, a potential solution to dangerous health misinformation might be rethinking the "endless" recommendation stream. Rather than suggesting any related video for any given issue, YouTube could curate sensitive issues such as health information, in the sense that only videos from certain verified accounts will be recommended to users. That said, it is worth noting that we found an antivaccination video created by a verified YouTube channel, highlighting just how difficult, but also important, this topic is.

Our research highlights the importance of conducting research on YouTube, not only to understand what is happening, but also to compel YouTube to be more transparent about its algorithms and how they influence how people watch news. Future research could go beyond the nonpersonalized platform-level recommendation approach chosen in our study and focus more on a personalized user-level approach. Although a more quantitative approach might not be feasible to analyze large-scale personalized YouTube recommendations, researchers could use qualitative methods, as they are commonly used in radicalization research (e.g., Baugut & Neumann, 2019), to find out more about the personal experience of users on YouTube as a platform and what role the algorithm plays. Furthermore, we believe in a more comparative perspective. Future research, for example, could focus on pandemics such as COVID-19 in an attempt to tease out the similarities and differences between COVID-19 and Zika. Many of the broader themes that we found for Zika—such as the origin of the virus, remedies, and the role of vaccinations—also seem to be present for COVID-19.

References

Allgaier, J. (2019). Science and environmental communication on YouTube: Strategically distorted communications in online videos on climate change and climate engineering. *Frontiers in Communication*, *4*, 1–15. https://doi.org/10.3389/fcomm.2019.00036

Allgaier, J. (2020). Science and medicine on YouTube. In J. Hunsinger, M. M. Allen, & L. Klastrup (Eds.), *Second international handbook of Internet research* (pp. 7–27). Springer VS. https://doi.org/10.1007/978-94-024-1555-1

Antony, M. G., & Thomas, R. J. (2010). 'This is citizen journalism at its finest': YouTube and the public sphere in the Oscar Grant shooting incident. *New Media & Society*, *12*(8), 1280–1296. https://doi.org/10.1177/1461444810362492

Baugut, P., & Neumann, K. (2019). Online propaganda use during Islamist radicalization. *Information, Communication & Society*, 1–23. https://doi.org/10.1080/1369118X.2019.1594333

Benkler, Y. (2006). *The wealth of networks: How social production transforms markets and freedom*. New Haven, CT: Yale University Press.

Benkler, Y., Faris, R., & Roberts, H. (2018). *Network propaganda: Manipulation, disinformation, and radicalization in American politics*. Oxford, UK: Oxford University Press.

Bora, K., Das, D., Barman, B., & Borah, P. (2018). Are Internet videos useful sources of information during global public health emergencies? A case study of YouTube videos during the 2015–16 Zika virus pandemic. *Pathogens and Global Health*, *112*(6), 320–328. https://doi.org/10.1080/20477724.2018.1507784

Bradshaw, S., Howard, P. N., Kollanyi, B., & Neudert, L.-M. (2020). Sourcing and automation of political news and information over social media in the United States, 2016–2018. *Political Communication*, *37*(2), 173–193. https://doi.org/10.1080/10584609.2019.1663322

Briones, R., Nan, X., Madden, K., & Waks, L. (2012). When vaccines go viral: An analysis of HPV vaccine coverage on YouTube. *Health Communication*, *27*(5), 478–485. https://doi.org/10.1080/10410236.2011.610258

Burgess, J., & Greenberg, J. (2009). *YouTube: Online video and participatory culture. Digital media and society series*. Cambridge, UK: Polity.

Carey, J. M., Chi, V., Flynn, D. J., Nyhan, B., & Zeitzoff, T. (2020). The effects of corrective information about disease epidemics and outbreaks: Evidence from Zika and yellow fever in Brazil. *Science Advances*, *6*(5), 1–10. https://doi.org/10.1126/sciadv.aaw7449

Chaves, M., & Braga, A. (2019). The agenda of disinformation: "Fake news" and membership categorization analysis in the 2018 Brazilian presidential elections. *Brazilian Journalism Research*, *15*(3), 474–495. https://doi.org/10.25200/BJR.v15n3.2019.1187

Covington, P., Adams, J., & Sargin, E. (2016). Deep neural networks for YouTube recommendations. In *RecSys '16: Proceedings of the 10th ACM Conference on Recommender Systems* (pp. 191–198). New York, NY: ACM. https://doi.org/10.1145/2959100.2959190

Dijck, J. van, & Poell, T. (2013). Understanding social media logic. *Media and Communication*, *1*(2), 2–14. https://doi.org/10.12924/mac2013.01010002

Egelhofer, J. L., & Lecheler, S. (2019). Fake news as a two-dimensional phenomenon: A framework and research agenda. *Annals of the International Communication Association*, *43*(2), 97–116. https://doi.org/10.1080/23808985.2019.1602782

Fisher, M., & Taub, A. (2019, August 11). How YouTube radicalized Brazil. *The New York Times*. Retrieved from https://www.nytimes.com/2019/08/11/world/americas/youtube-brazil.html

Freelon, D., & Wells, C. (2020). Disinformation as political communication. *Political Communication*, *37*(2), 145–156. https://doi.org/10.1080/10584609.2020.1723755

Gabarron, E., Fernandez-Luque, L., Armayones, M., & Lau, A. Y. (2013). Identifying measures used for assessing quality of YouTube videos with patient health information: A review of current literature. *Interactive Journal of Medical Research*, *2*(1), 1–9. https://doi.org/10.2196/ijmr.2465

Golebiewski, M., & boyd, d. (2018, May). Data voids: Where missing data can easily be exploited. *Data & Society*. https://datasociety.net/wp-content/uploads/2018/05/Data_Society_Data_Voids_Final_3.pdf

Guess, A., Nagler, J., & Tucker, J. (2019). Less than you think: Prevalence and predictors of fake news dissemination on Facebook. *Science Advances*, *5*(1), 1–9. doi:10.1126/sciadv.aau4586

Guess, A. M., Nyhan, B., & Reifler, J. (2018). *Selective exposure to misinformation: Evidence from the consumption of fake news during the 2016 U.S. presidential campaign*. Retrieved from http://www.ask-force.org/web/Fundamentalists/Guess-Selective-Exposure-to-Misinformation-Evidence-Presidential-Campaign-2018.pdf

Guo, L., & Vargo, C. (2020). "Fake news" and emerging online media ecosystem: An integrated intermedia agenda-setting analysis of the 2016 U.S. presidential election. *Communication Research*, *47*(2), 178–200. https://doi.org/10.1177/0093650218777177

Jacobs, A. (2019, July 2). The Zika virus is still a threat. Here's what the experts know. *The New York Times*. Retrieved from https://www.nytimes.com/2019/07/02/health/zika-virus.html

Kaiser, J., & Rauchfleisch, A. (forthcoming). Birds of a feather get recommended together: Algorithmic Homophily in YouTube's channel recommendations in the United States and Germany. *Social Media + Society*.

Kaiser, J., & Rauchfleisch, A. (2018, April 11). Unite the right? How YouTube's recommendation algorithm connects the U.S. far-right. *Data & Society*. Retrieved from https://medium.com/@MediaManipulation/unite-the-right-how-youtubes-recommendation-algorithm-connects-the-u-s-far-right-9f1387ccfabd

Kaiser, J., Rauchfleisch, A., & Bourassa, N. (2020). Connecting the (far-)right dots: A topic modeling and hyperlink analysis of (far-)right media coverage during the U.S. elections 2016. *Digital Journalism*, 8(3), 422–441. https://doi.org/10.1080/21670811.2019.1682629

Kim, J. (2012). The institutionalization of YouTube: From user-generated content to professionally generated content. *Media, Culture & Society, 34*(1), 53–67. https://doi.org/10.1177/0163443711427199

Klofstad, C. A., Uscinski, J. E., Connolly, J. M., & West, J. P. (2019). What drives people to believe in Zika conspiracy theories? *Palgrave Communications*, *5*(1), 1–8. https://doi.org/10.1057/s41599-019-0243-8

Lewandowsky, S., Ecker, U. K. H., & Cook, J. (2017). Beyond misinformation: Understanding and coping with the "post-truth" era. *Journal of Applied Research in Memory and Cognition*, *6*(4), 353–369. https://doi.org/10.1016/j.jarmac.2017.07.008

Lowe, R., Barcellos, C., Brasil, P., Cruz, O., Honório, N., Kuper, H., & Carvalho, M. (2018). The Zika virus epidemic in Brazil: From discovery to future implications. *International Journal of Environmental Research and Public Health*, *15*(1), 1–18. https://doi.org/10.3390/ijerph15010096

Nerghes, A., Kerkhof, P., & Hellsten, I. (2018). Early public responses to the Zika-virus on YouTube. In *WebSci '18*: *Proceedings of the 10th ACM Conference on Web Science* (pp. 127–134). New York, NY: ACM. https://doi.org/10.1145/3201064.3201086

Newman, N., Fletcher, R., Kalogeropoulos, A., & Nielsen, R. K. (2019). *Reuters Institute Digital News Report 2019*. Oxford, UK: Reuters Institute.

O'Callaghan, D., Greene, D., Conway, M., Carthy, J., & Cunningham, P. (2015). Down the (white) rabbit hole: The extreme right and online recommender systems. *Social Science Computer Review*, *33*(4), 459–478. https://doi.org/10.1177/0894439314555329

Pant, S., Deshmukh, A., Murugiah, K., Kumar, G., Sachdeva, R., & Mehta, J. L. (2012). Assessing the credibility of the "YouTube approach" to health information on acute myocardial infarction. *Clinical Cardiology*, *35*(5), 281–285. https://doi.org/10.1002/clc.21981

Rauchfleisch, A., & Kaiser, J. (2020). The German far-right on YouTube: An analysis of user overlap and user comments. *Journal of Broadcasting & Electronic Media*, 1–24. https://doi.org/10.1080/08838151.2020.1799690

Reinhart, R. (2018). *Brazilians face confidence crisis ahead of election*. Retrieved from https://news.gallup.com/poll/243161/brazilians-face-confidence-crisis-ahead-key-election.aspx

Roberts, M., Stewart, B., & Tingley, D. (2019). stm: An R package for structural topic models. *Journal of Statistical Software, 91*(2), 1–40. http://dx.doi.org/10.18637/jss.v091.i02

Roth, C., Mazières, A., & Menezes, T. (2020). Tubes and bubbles topological confinement of YouTube recommendations. *PLOS ONE*, *15*(4), 1–17. https://doi.org/10.1371/journal.pone.0231703

Sood, A., Sarangi, S., Pandey, A., & Murugiah, K. (2011). YouTube as a source of information on kidney stone disease. *Urology*, *77*(3), 558–562. https://doi.org/10.1016/j.urology.2010.07.536

Syed-Abdul, S., Fernandez-Luque, L., Jian, W.-S., Li, Y.-C., Crain, S., Hsu, M.-H., . . . & Liou, D.-M. (2013). Misleading health-related information promoted through video-based social media: Anorexia on YouTube. *Journal of Medical Internet Research*, *15*(2), 1–13. https://doi.org/10.2196/jmir.2237

Traag, V. A. (2015). Faster unfolding of communities: Speeding up the Louvain algorithm. *Physical Review E*, *92*(3), 1–10. https://doi.org/10.1103/PhysRevE.92.032801

Tucker, J., Guess, A., Barbera, P., Vaccari, C., Siegel, A., Sanovich, S., . . . & Nyhan, B. (2018). Social media, political polarization, and political disinformation: A review of the scientific literature. *SSRN Electronic Journal*, 1–95. https://doi.org/10.2139/ssrn.3144139

Wardle, C. (2018). The need for smarter definitions and practical, timely empirical research on information disorder. *Digital Journalism*, *6*(8), 951–963. https://doi.org/10.1080/21670811.2018.1502047

Wissenschaft im Dialog. (2018). *Tabellenband Wissenschaftsbarometer 2018* [Science barometer book of tables 2018]. Retrieved from https://www.wissenschaft-im-dialog.de/fileadmin/user_upload/Projekte/Wissenschaftsbarometer/Dokumente_18/Downloads_allgemein/Tabellenband_Wissenschaftsbarometer2018_final.pdf

YouTube. (2020). *YouTube for Press*. Retrieved from https://www.youtube.com/about/press/

Zhao, Z., Chi, E., Hong, L., Wei, L., Chen, J., Nath, A., . . . & Yi, X. (2019). Recommending what video to watch next: A multitask ranking system. In *RecSys '19: Proceedings of the 13th ACM Conference on Recommender Systems* (pp. 43–51). New York, NY: ACM. https://doi.org/10.1145/3298689.3346997

Zimmermann, F., & Kohring, M. (2020). Mistrust, disinforming news, and vote choice: A panel survey on the origins and consequences of believing disinformation in the 2017 German parliamentary election. *Political Communication*, *37*(2), 215–237. https://doi.org/10.1080/10584609.2019.1686095