

Multilingual Analysis of YouTube’s Recommendation System: Examining Topic and Emotion Drift in the ‘Cheng Ho’ Narrative

Ugochukwu Onyepunuka¹, Mustafa Alassad¹, Lotenna Nwana¹ and Nitin Agarwal¹

¹COSMOS Research Center, University of Arkansas - Little Rock, USA

Abstract

YouTube is a major source of information for many users and its recommendation algorithm is pivotal in video discovery and viewership on the platform. It is responsible for 70% of the content users engage with on its platform, and the type of information users are exposed to. The importance of scrutinizing recommendation systems, to understand how potential algorithmic bias may affect users in spreading disinformation, cannot be overemphasized. Previous studies have shown that the recommendation algorithm has an inherent tendency of bias toward a small fraction of videos, and it pushes users into mild ideological echo chambers. This study aims to determine the extent to which YouTube’s recommendation algorithm spreads disinformation, using the Cheng Ho narrative. Cheng Ho was a Muslim Chinese naval admiral in the 15th century, nicknamed the “Chinese Columbus”. He was a symbol of China’s Islamic diplomacy and peaceful ascendancy to power. To achieve the study’s aim, a list of 50 videos on Cheng Ho was collected by passing relevant keywords to YouTube’s search API. These 50 videos served as the seed for the recommendations, with 58,825 unique videos collected through 5 depths of recommendations. We computed the topic drift on the recommendation depths and discovered that the recommendations led us further away from the original topic. Furthermore, observing the eigenvector centrality values of videos within the recommendation network of different depths, we saw the evolution of influential videos as their relevance to Cheng Ho diminished. The results showed how YouTube’s recommendation system discards the topics of the seed videos by subtly introducing a new but still pro-China topic in the network through influential videos. This new topic is about economic growth and religious freedom in China targeting Indonesia’s younger demographic by focusing on current events and pop culture. This study sets the stage for further research in analyzing bias in recommendation algorithms, their exploitation by information actors, their impact on mis/disinformation propagation, and their effect on user consumption.

Keywords

Cheng Ho, China, disinformation, recommendation bias, topic drift, YouTube

1. Introduction

YouTube is the largest video-sharing platform today, with over 6 billion hours of videos watched by its visitors every month [1]. It has become a major source of information for most people, as over 2.5 billion active users use it monthly [2]. The content on YouTube covers a vast range of

In: R. Campos, A. Jorge, A. Jatowt, S. Bhatia, M. Litvak (eds.): Proceedings of the Text2Story’23 Workshop, Dublin (Republic of Ireland), 2-April-2023

*Corresponding author.

✉ uonyepunuka@ualr.edu (U. Onyepunuka); mmalassad@ualr.edu (M. Alassad); ltnwana@ualr.edu (L. Nwana); nxagarwal@ualr.edu (N. Agarwal)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

categories, including education, sports, politics, and religion. YouTube recommendation was designed to suggest content based on a user's current and previous viewings. The recommendation algorithm on the platform is intended to maximize user retention by suggesting content appealing to a user's interest. With YouTube being a critical source of information for many users and its recommendation algorithm driving the type of information users are exposed to, it is important to scrutinize the recommendation system to understand how potential algorithmic bias may affect users. Consequences of algorithmic bias, such as misinformation and disinformation could have great impacts on discrimination against marginalized communities. For this study, we analyze the algorithm for the presence of bias, and if the detected bias aids in the spread of the Cheng Ho propaganda narrative.

Cheng Ho, also known as Zheng He, was a Chinese naval admiral who commanded naval voyages in the early 15th century through Southeast Asia, India, and the Middle East [3]. The Cheng Ho mythology seeks to advance China's Islamic diplomacy by portraying him as a benevolent giver who spread Islam across Southeast Asia [3] [4] [5]. The manipulation of the Cheng Ho narrative is intended to increase regional support for China's "Maritime Silk Road" [5][6], and as a response to their oppression of Uyghur Muslims [7]. Accompanied by the public scrutiny of China's oppression of the Uyghur Muslims, the Chinese Communist Party (CCP) has invested in reviving and manipulating the Cheng Ho myth [3] [4]. The Cheng Ho myth has been weaved into a symbol of China's peaceful ascendancy to power while portraying China's economic, military, and naval prowess [5] [6]. Peggy-Jean et al. [3] highlight the importance of this disinformation campaign to China's geopolitical ambitions in the South China Sea. The tactics deployed by these disinformation actors rely on constructing collective memory by repackaging history in a fictional and false manner [3]. Furthermore, Geoff Wade et al. [6] present a revisionist view of Cheng Ho's voyages and argue that the voyages of friendships were an aggressive attempt to establish China's dominance in places like Vietnam.

The motivation for this study comes from the need to highlight the extent of algorithmic bias in spreading disinformation. Through this research, we found a set of highly-influential YouTube videos that acted as attractors in the recommendation network. In a further analysis, we examined the characteristics of these attractors and determined their relevance to the Cheng Ho narrative. A multilingual analysis of the video recommendations was also performed to compare patterns that may exist in different cultural contexts.

The results showed that content related to our seed videos was filtered out across recommendations. Simultaneously, new content unrelated to the seed videos were introduced into the network through attractors. These new content identified as 'pro-China' topics in the network, which focused on economic growth and religious freedom in China, also targeted Indonesia's younger demographic by incorporating current events and pop culture.

2. Literature Review

This section presents related works on, topic drift, recommendation bias, disinformation, and radicalization. In recent years the algorithmic bias of recommendation engines has been studied to understand the extent of their contribution to spreading misinformation, leading users into echo chambers, polarization, and radicalization.

Disinformation is intentional false information designed to deceive or mislead people [8]. Undoubtedly, political disinformation narratives are intended to influence people's perceptions of reality to advance authoritarian and populist agendas [9]. Additionally, it aims to inflame polarization and discrimination against marginalized communities, subvert human rights defenders and human rights processes, and discredit facts [9]. Despite its negative impact on society, little research has been conducted to understand its heterogeneous effect on minority groups. Inspired by this, Neo et al. [10] analyzed the qualitative effects of disinformation on Indonesia's racial, ethnic, and sexual minority communities. The study was conducted on data from interviews with Indonesian citizens belonging to minority groups. It revealed how dominant social groups had utilized disinformation as a tool to gather various types of political and religious capital and socially control and punish minority groups. Furthermore, they concluded that given the spread of fake news through social media platforms, increased regulation of these platforms has the capacity to mitigate the effects of disinformation and foster a healthier society [10].

Brown et al. [11] analyzed the extent to which YouTube pushes users into rabbit holes or echo chambers of ideologically biased content. To achieve this, they developed a method to estimate the ideology of videos and avoided user personalization from the recommendation algorithm. They discovered that the recommendation algorithm pushes users into mild echo chambers, but insufficient evidence indicates that it leads them down the rabbit hole of ideologically extreme content [11]. In conclusion, they observed that the longer users follow the recommendations they are pulled into a narrow range of ideological content regardless of the users' ideology. Kirdemir et al. [12] described similar findings showing YouTube's recommendation bias in favor of a streamlined set of contents under generalizable conditions.

Heuer et al. [13] research aim was to audit the bias that emanates from YouTube's recommendation algorithm on political topics in Germany. They followed ten chains of recommendations per video on the political topic to examine potential recommendation bias from YouTube. Their findings suggest that YouTube enacts strong popularity bias in its recommendations, but the videos are topically dissimilar or unrelated to the original narrative [13]. The researchers also discovered emotion drift from sadness to happiness in the recommendations and examined the relationship between content popularity and emotions.

Other related works have employed the use of human annotators in content analysis to quantify topic and emotion drift. This study presents a systematic computational approach including Topic Modeling [14], Network Analysis [15], and Hellinger Distance Score [16] to examine the extent of YouTube's recommendation bias in spreading a disinformation narrative. Unlike other works, this study provides a multilingual analysis of the recommendation algorithm to compare patterns in different cultural contexts. Through this study, we examined the evolution of a YouTube recommendation network on a disinformation narrative (Cheng Ho) as it progresses through the recommendation depths.

3. Methodology

This section describes the data collection process and the methodology utilized in this study.

Table 1
Recommendation Depth Statistics.

Depth	Number of Videos	Number of Unique Videos
Seed	50	50
Depth 1	247	188
Depth 2	1212	819
Depth 3	5985	3521
Depth 4	29586	14755
Depth 5	145923	47101

3.1. Data Collection

YouTube API crawler described by Kready et al. [17] was used to collect over 1000 videos based on the following keyphrases:

“Cheng Ho’/Zheng He’+laksamana+damai”, “Sam Po Kong’+Islam+Indonesia”, “1421 Saat China Menemukan Dunia’+Gavin Menzies’”, and “1421 Saat China Menemukan Dunia”.

The data collected was written to a MySQL database, where the video titles were queried with a full-text search of the keyphrases. Furthermore, we extracted the top 50 most viewed videos to limit the results.

Similarly, we adopted a custom crawler to collect five depersonalized video recommendations on each seed video, producing the 1st depth of recommendations. Furthermore, this process was repeated four more times, with videos from the previous depths used as inputs to generate the next depths. The depersonalization was done to evaluate the raw recommendation algorithm without the influence of the user history. Table 1. highlights the number of videos collected at each depth.

3.2. Topic Drift

By analyzing the topic drift, we aimed to determine if recommendations deviate from the original topic and investigate the content similarity that occurs in YouTube’s recommendation system as it progresses through five depths of recommendations. Topic drift occurs when the topics of the recommended videos deviate from the original topics found in the seed videos. To measure topic drift, the titles of unique videos across all depths were concatenated to create a master corpus. The master corpus was subjected to an LDA topic model [14] to determine the topics present, which revealed 10 topics in the master corpus. Next, a corpus was created for each depth by combining the titles of all the videos in a depth. Furthermore, the corpus for each depth was passed through the LDA topic model to generate their topic probability distribution (likelihood for a corpus to belong to a topic). Finally, to measure content similarity between two depths we applied the Hellinger Distance metric [16] to their topic probability distribution. The Hellinger distance formula is a mathematical formula used to measure the similarity or dissimilarity between two probability distributions. Given two probability distributions P and Q, the Hellinger distance $H(P, Q)$ is defined as:

$$H(P, Q) = \frac{1}{\sqrt{2}} \sqrt{\sum_{i=1}^k (\sqrt{p_i} - \sqrt{q_i})^2} \quad (1)$$

The Hellinger distance ranges from 0 (when P and Q are identical) to 1 (when P and Q have no common outcomes). The topic drifts were measured in two ways:

- Seed-to-depth distance: It quantifies the topic similarity between the seed videos (depth 0) and each depth (e.g depth 0 & depth 1; depth 0 & depth 2; depth 0 & depth 3) by measuring the distance between their topic probability distribution.
- Inter-depth distance: It quantifies the topic similarity between adjacent depths (e.g. depth 0 & depth 1; depth 1 & depth 2; depth 2 & depth 3) by measuring the distance between their topic probability distribution.

3.3. Network Analysis

We performed a network analysis to understand the network structure within the recommendation depths. The aim was to observe how the recommendation network evolves as we progress through the recommendation depths. To analyze the recommendation network, we created a network graph for each depth and analyzed it using Gephi [18]. The recommendation network consists of videos as vertices and recommendations as directed edges. Next, we used the modularity measure [19] on Gephi to identify communities within the network, which enabled us to isolate the videos in a community and summarize the videos' characteristics to understand their uniqueness. We utilized the eigenvector centrality value (eigenvalues) to measure the transitive influence of nodes in the network [20]. The value (0 to 1) of a node is dependent on how well-connected a node is to other well-connected nodes. Applying the eigenvector centrality measure to the network helped identify the influential videos (high eigenvalues) acting as attractors in the recommendation network. These videos with high eigenvalues are favored by the algorithm with a propensity to be recommended. Furthermore, we determined the similarity of the influential videos to the original narrative.

3.4. Emotion Analysis

Next, we performed an emotional analysis to examine the emotions embedded in the recommendations and measure emotion drift (the divergence of emotions from the original narrative). This allowed us to further assess how the Youtube recommendation algorithm considers emotions [13]. Emotion analysis involves detecting the true emotions behind a text. For effective text analysis, the emotion analysis technique employs natural language processing, information extraction techniques, and a pre-trained language representation model [21]. Manuel Romero's T5-based-finetuned-emotion model [22] was used to identify the emotions of video titles, using transfer learning which preserves previous information, thereby improving the speed and accuracy of training. This produced a probability score for six emotions: anger, fear, joy, love, surprise, and sadness. Furthermore, we computed the distribution of emotions for each depth and visualized it on a line graph, allowing us to see the emotion trends as we traverse through the recommendation depths. This process was conducted on the subset of Indonesian and

English video titles for each depth, converting the Indonesian video title to English before passing it as input into the emotion model.

4. Results

This section discusses our findings at each stage of our methodology.

4.1. Topic Drift Analysis of Recommended Videos

The topic drift analysis aims to determine if YouTube's recommendations deviate from the original topic as it progresses through five depths of recommendations. A high Hellinger distance score signifies low similarity between two probability distributions and vice versa. Therefore, topic drift occurs when there is a sequential increase in the Hellinger distance between two topic probability distributions. From the line chart, an upward trend indicates topic drift (increasing Hellinger distance) explaining that there is decreasing similarity.

After collecting the recommendations for five depths we calculated the language distribution of videos in each depth. Consequently, we saw that Indonesian and English had the largest percentage, 55.5%, and 34.6% respectively, thereby influencing our decision to perform a comparative analysis of the two languages. The topic drift methodology described above was applied to a subset of these languages across depths. Since our seed videos were all Indonesian, we translated them into English and used them to calculate the topic drift of the English videos

In Figure 1 (a & b), we measured the topic similarity of the recommended videos between adjacent depths (the inter-depth distance). Looking at the results of the inter-depth distance of English and Indonesian videos in Figure 1(a & b), we observed a declining Hellinger distance score. From the graphs, the maximum distance occurred between the seed (0) and depth 1, with a constant decline to the end. This explains that as we progress through the recommendation depths the content similarity between two adjacent depths increases. It tells us that the recommendation becomes more similar to each other.

Furthermore, the seed-to-depth Hellinger distance score for English and Indonesian videos in Figure 1 (c & d), was used to determine the relevance of the recommendations in each depth to the original narrative. From the graphs, we observed that the distance between the seed and the recommendations increases as we progress through the depths. The observed continuous topic drift explains that the recommended videos were becoming more dissimilar from the seed as we moved through the depths.

Combining the results of the two topic drift views, suggests that regardless of each depth's increasing dissimilarity to the seed (seed-depth distance), YouTube's recommendation engine strives to keep the content between each adjacent depth similar (inter-depth distance). This explains how users are gradually exposed to videos irrelevant to the seed. We realized a higher magnitude in topic drifts for English videos over Indonesian videos.

4.2. Analysis of the Recommendation Network

Next, we analyzed each depth's network to identify the influential nodes with a higher propensity to be recommended by the algorithm. The videos in each depth were ranked with their




Figure 1: Topic Drifts (a)Inter-depth distance, English videos (b)Inter-depth distance, Indonesian videos (c)Seed-depth distance, English videos (d)Seed-depth distance, Indonesian videos.

eigenvector centrality value, which determines a node’s transitive influence in the network. We selected and examined the top 3 most influential videos based on the eigenvalues. By examining the network’s influential nodes, we confirm the kinds of videos that affect recommendations and how a video’s influence evolves as we go through the depths. We observed how the influential nodes in the network evolve from relevant videos to irrelevant videos about the Cheng Ho topic. The topics of the influential nodes drift from the original narrative and a new topic (“Cha Guan”) is subtly introduced at depths 2 and 3, with a complete shift at depth 4. “Cha Guan” is a segment on the “Asumi” channel, a media-tech institution aimed at Indonesia’s younger demographic, with a focus on current events and pop culture. Notably, the videos on this channel discussed the economy and religious freedom in China. Still observing the evolution of influential videos in the network, we noticed a shift in language at depth 5, from Indonesian to English. In addition to the language change amongst influential videos at depth 5, the distribution of eigenvector scores was significantly lower.

4.3. Emotion Drift Analysis

Analyzing the emotions attached to videos at different recommendation depths and whether they remain similar to the original narrative aids to assess the impact of emotions on YouTube’s recommendations. Performing a comparative analysis of the emotion drift on Indonesian and English videos across all depths, allows us to identify patterns that may exist in the different language contexts. However, in the emotion drift charts of Indonesian and English videos shown in Figure 5, we see that a similar trend exists between them. Joy is the most prominent emotion exhibited in the video titles of the different languages. This could be due to the likelihood for users to engage more with positive emotions than negative ones [13] or that it aligns with the disinformation propaganda of the Cheng Ho narrative driven by China. Further analysis of the impact of emotions on video recommendations is conducted in future work by comparing the emotion drifts on competing narratives.

The results from the topic drift analysis show high content similarity in the recommendations between depths, but low content similarity with the seeds. Accordingly, the increased similarity between recommendation depths and their divergence from the seed videos is indicative of bias toward a coordinated pool of videos with no relevance to the original topic. Furthermore, analyzing the network with the eigenvector centrality measure revealed influential videos at each depth. Additionally, the emotion analysis performed on the data showed joy as the most

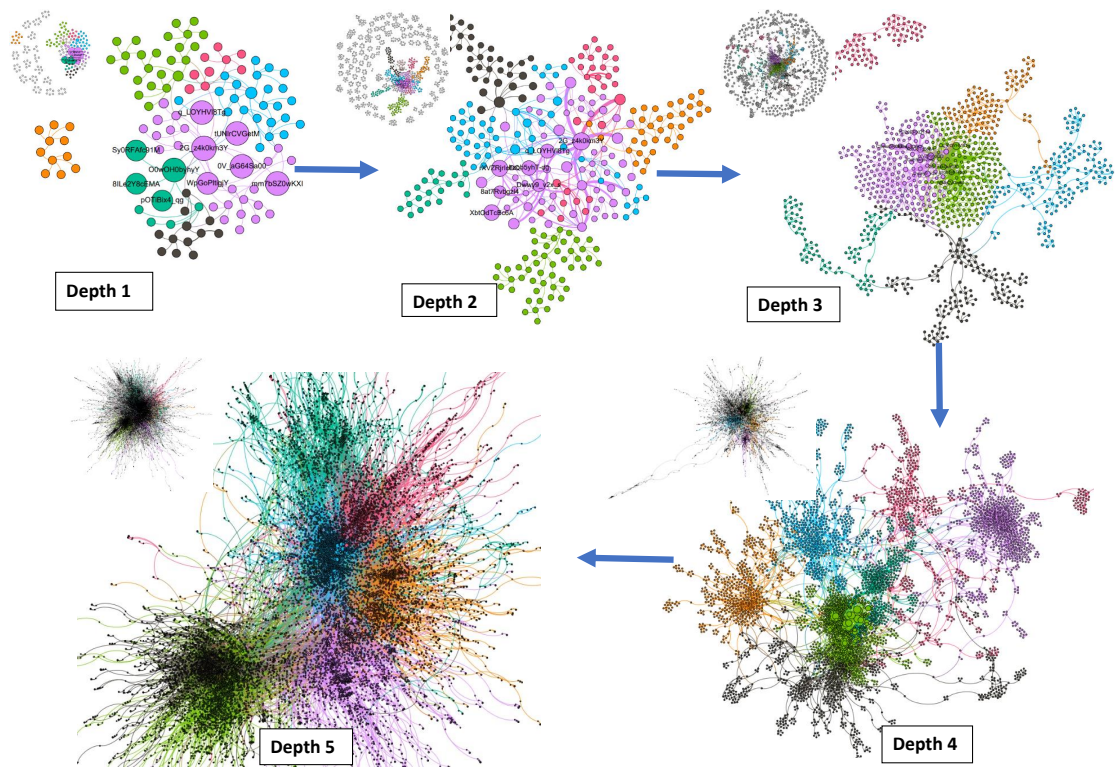


Figure 2: Evolution of YouTube’s recommendation network.

prominent emotion, which aligns with emotions attached to the Cheng Ho narrative.

5. Conclusion

From our analysis, we find no evidence affirming the contributions of YouTube’s recommendations toward spreading disinformation. The results of our topic drift analysis show that as we progressively step through YouTube’s recommendations, we get further away from our original narrative. However, we observed an increased similarity between the adjacent recommendation depths. The increased similarity between recommendations and their divergence from the seed videos is suggestive of a bias towards a streamlined pool of videos with no relevance to the original topic. This is evident in the results of our network analysis, where we see how the content of the influential videos switches from our original topic to videos about “Cha Guan”. From our comparison of Indonesian and English topic drifts, we observed similar patterns between them, but the English videos had no relevance to our original topic. The emotional analysis conducted on the recommended videos showed joy as the most prominent emotion, this could either be due to the likelihood for users to engage more with positive emotions

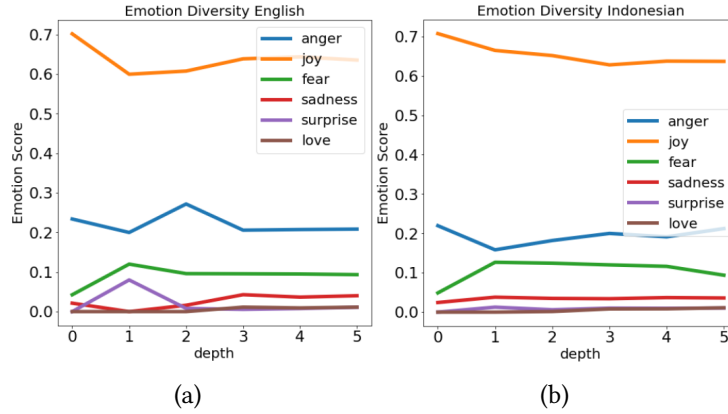


Figure 3: Emotion drift analysis of (a) English Videos and (b) Indonesian Videos.

than negative ones or tie into the theme of the China-driven Cheng Ho narrative. From our findings, we conclude that YouTube’s recommendation algorithm does not aid the spread of disinformation, but that as we step through recommendations the relevance of their contents to the original narrative is reduced.

Future directions for this research include evaluating topic drifts on different levels of video text including video description, transcript, and a concatenation of video title and description. Other directions may involve building a predictive model from the video characteristics (engagement stats, eigenvector score, topic distribution, etc.) at each depth to determine the likelihood of a video being recommended by another video.

Acknowledgments

This research is funded in part by the U.S. National Science Foundation (OIA-1946391, OIA-1920920, IIS-1636933, ACI-1429160, and IIS-1110868), U.S. Office of the Under Secretary of Defense for Research and Engineering (FA9550-22-1-0332), U.S. Office of Naval Research (N00014-10-1-0091, N00014-14-1-0489, N00014-15-P-1187, N00014-16-1-2016, N00014-16-1-2412, N00014-17-1-2675, N00014-17-1-2605, N68335-19-C-0359, N00014-19-1-2336, N68335-20-C-0540, N00014-21-1-2121, N00014-21-1-2765, N00014-22-1-2318), U.S. Air Force Research Laboratory, U.S. Army Research Office (W911NF-20-1-0262, W911NF-16-1-0189, W911NF-23-1-0011), U.S. Defense Advanced Research Projects Agency (W31P4Q-17-C-0059), Arkansas Research Alliance, the Jerry L. Maulden/Entergy Endowment at the University of Arkansas at Little Rock, and the Australian Department of Defense Strategic Policy Grants Program (SPGP) (award number: 2020-106-094). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding organizations. The researchers gratefully acknowledge the support.

References

- [1] YouTube: What is YouTube?, 2022. URL: <https://edu.gcfglobal.org/en/youtube/what-is-youtube/1/>.
- [2] Biggest social media platforms 2022, 2022. URL: <https://www.statista.com/statistics/272014/global-social-networks-ranked-by-number-of-users/>.
- [3] P.-J. Allin, S. Corman, “China’s Columbus” Was an Imperialist Too: Contesting the Myth of Zheng He (2022). URL: <https://smallwarsjournal.com/jrnl/art/chinas-columbus-was-imperialist-too-contesting-myth-zheng-he>.
- [4] ‘Fake news of the century’: The Muslim explorer China deploys while persecuting Muslims, ABC News (2019). URL: <https://www.abc.net.au/news/2019-09-22/zheng-he-chinese-islam-explorer-belt-and-road/11471758>.
- [5] Zheng He and the maritime silk road, 2015. URL: <https://u.osu.edu/mclc/2015/10/02/zheng-he-and-the-maritime-silk-road/>.
- [6] G. Wade, The Zheng He Voyages: A Reassessment, *Journal of the Malaysian Branch of the Royal Asiatic Society* 78 (2005) 37–58. doi:10.2307/41493537.
- [7] China’s Repression of Uyghurs in Xinjiang, 2022. URL: <https://www.cfr.org/backgrounder/china-xinjiang-uyghurs-muslims-repression-genocide-human-rights>.
- [8] Misinformation and disinformation, 2022. URL: <https://www.apa.org/topics/journalism-facts/misinformation-disinformation>.
- [9] APC, Disinformation and freedom of expression: Submission in response to the call by the UN Special Rapporteur on the promotion and protection of the right to freedom of opinion and expression | Association for Progressive Communications, 2021. URL: <http://bit.ly/3FPHghD>.
- [10] A. A. Ric Neo, Jason DC Yin, Of social discipline and control: The impact of fake news and disinformation on minorities in Indonesia | Association for Progressive Communications, 2021. URL: <https://www.apc.org/en/pubs/social-discipline-and-control-impact-fake-news-and-disinformation-minorities-indonesia>.
- [11] M. A. Brown, J. Bisbee, A. Lai, R. Bonneau, J. Nagler, J. A. Tucker, Echo Chambers, Rabbit Holes, and Algorithmic Bias: How YouTube Recommends Content to Real Users, 2022. URL: <https://papers.ssrn.com/abstract=4114905>. doi:10.2139/ssrn.4114905.
- [12] B. Kirdemir, J. Kready, E. Mead, M. Hussain, N. Agarwal, Examining Video Recommendation Bias on YouTube, 2021, pp. 106–116. doi:10.1007/978-3-030-78818-6_10.
- [13] H. Heuer, H. Hoch, A. Breiter, Y. Theocharis, Auditing the Biases Enacted by YouTube for Political Topics in Germany, in: *Mensch und Computer 2021*, 2021, pp. 456–468. URL: <http://arxiv.org/abs/2107.09922>. doi:10.1145/3473856.3473864, arXiv:2107.09922 [cs].
- [14] Topic Modeling and Latent Dirichlet Allocation (LDA) | DataScience+, 2022. URL: <https://datascienceplus.com/topic-modeling-and-latent-dirichlet-allocation-lda/>.
- [15] Network Analysis - an overview | ScienceDirect Topics, 2018. URL: <https://www.sciencedirect.com/topics/social-sciences/network-analysis>.
- [16] A. Shemyakin, Hellinger Distance and Non-informative Priors, *Bayesian Analysis* 9 (2014) 923–938. URL: <https://projecteuclid.org/journals/bayesian-analysis/volume-9/issue-4/Hellinger-Distance-and-Non-informative-Priors/10.1214/14-BA881.full>. doi:10.1214/14-BA881, publisher: International Society for Bayesian Analysis.

- [17] J. Kready, S. A. Shimray, M. N. Hussain, N. Agarwal, YouTube Data Collection Using Parallel Processing, in: 2020 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), 2020, pp. 1119–1122. doi:10.1109/IPDPSW50202.2020.00185.
- [18] Learn how to use Gephi, 2008. URL: <https://gephi.org/users/>.
- [19] Data Story: Gephi - Clustering layout by modularity, 2014. URL: <https://parklize.blogspot.com/2014/12/gephi-clustering-layout-by-modularity.html>.
- [20] Eigenvector Centrality - Neo4j Graph Data Science, 2023. URL: <https://neo4j.com/docs/graph-data-science/2.2/algorithms/eigenvector-centrality/>.
- [21] N. Shelke, S. Chaudhury, S. Chakrabarti, S. L. Bangare, G. Yogapriya, P. Pandey, An efficient way of text-based emotion analysis from social media using LRA-DNN, *Neuroscience Informatics* 2 (2022) 100048. URL: <https://www.sciencedirect.com/science/article/pii/S2772528622000103>. doi:10.1016/j.neuri.2022.100048.
- [22] mrm8488/t5-base-finetuned-emotion · Hugging Face, 2022. URL: <https://huggingface.co/mrm8488/t5-base-finetuned-emotion>.