Examining Content and Emotion Bias in YouTube's Recommendation Algorithm

Obianuju Okeke, Mert Can Cakmak, Billy Spann, Nitin Agarwal
COSMOS Research Center
University of Arkansas at Little Rock (UALR)
Little Rock, USA
Email: {oiokeke, mccakmak, bxspann, nxagarwal}@ualr.edu

***Abstract*—Detection, characterization, and mitigation of bias in modern systems of automated and autonomous decisions is a growing interdisciplinary field. This study aims to explore YouTube's video recommendation bias to determine if an inherent bias has an unintended impact of occluding vulnerable communities and minority groups. Our findings suggest that the algorithm recommends videos evoking more positive emotions and higher user engagement. We also discovered that content related to our seed videos was filtered out in a systematic but gradual pendulum-like motion. This analysis of potential emergent biases will be applicable in analyzing the fairness of recommender systems, patterns of content consumption, information diffusion, echo-chamber formation, and other significant problems.**

***Index Terms*—Keywords-Recommender Systems; Recommendation Bias; YouTube; Topic Modeling; Emotion Modeling.**

I. Introduction

According to YouTube's Chief Product Officer Neal Mohan [3], around 70 percent of videos watched on YouTube are recommended videos, this means that an average of 7 out of 10 videos a user watches are recommended by YouTube. Although YouTube's goal of profit generation through increased watch-time is intended to be harmless and business-oriented, this pattern could have the unintended consequence of occluding vulnerable communities and crisis-torn societies. For our research, we studied the impact of the algorithm on videos related to the Uyghur group, a vulnerable community in the China-Uyghur crisis. According to the Council on Foreign Relations, more than a million Uyghurs - a Muslim, Turkish speaking ethnic group, have been detained since 2017 in the China Xinjiang region [15]. Platforms such as YouTube remain an indispensable outlet for such minority groups and vulnerable communities to spread awareness on important issues [21]. It also serves as a window to the world to receive vital information [19]. These groups depend on free and open platforms such as YouTube to vocalize the crisis they endure in their respective societies. According to Silverman, content evoking polarization is propagated faster than non-polarizing content [22]. We, therefore, expect content and emotions related to our seed videos to be propagated across recommendation depths.

II. Literature Review

In this section, we discuss research related to our study which includes previous works on topic shifting, emotion detection [1], and bias in recommender systems. Bias in recommendation engines has been extensively studied to understand its nature, structure, and effects, especially in the area of radicalization, polarization, and spread of misinformation [18]. These studies have described how homophilic communities are generated through recommended videos as well as factors which drive the creation of such interconnected communities, leading to filter bubble effects and echo-chambers [23]. Insights from such studies are crucial in identifying the emergence of homogeneity in recommender systems. Topic drift is a technique that has been used by many researchers in studying how content evolves. By studying content evolution, we are able to determine if content remains the same or changes relative to a standard metric. O' Hare et al. [7] analyzed sentiment-annotated corpus of textual data to determine topic drift among documents within a corpus. Liu et al. developed an LDA (Latent Dirichlet Allocation)-based method for topic drift detection in micro-blog posts [5] Topal et al. identified and quantitatively studied the effects of topic shift in social media comments [17]. Papakyriakopoulos et al., addressed hyperactive users and their effects on political discussion and recommender systems [13]. According to Papakyriakopoulos, recommendation algorithms favor the interest of hyperactive users, creating significant social influence bias and causing alterations in political opinions. By identifying inherent topics using topic modeling [4], [9], the authors classified content by topic to examine the activities of hyperactive users and determine if engagement distribution diverges. In this paper, we aim to identify inherent bias in YouTube's recommendation algorithm, and determine if the identified bias works to occlude videos related to vulnerable communities across recommendation depths. Some of the questions we hope to answer include:

- **RQ1**: How do we identify bias in content related to vulnerable communities?
- **RQ2**: What kind of videos drive recommendations on YouTube?
- **RQ3**: How do videos related to vulnerable communities change across recommendation depth?

Unlike other methodologies which have adopted a more manual approach through the use of raters in the content analysis process [16], we programmatically assign topic communities to videos across recommendation depths. Through our research, we also aim to track the evolution of content across recommendations for a detailed view on how content related to the Uyghur ethnic group is recommended on YouTube. In the next section, we present our data collection methodology.

III. Data Collection

To begin data collection, we conducted expert workshops to identify keywords related to the China-Uyghur issue. These keywords were used as search queries on YouTube's search engine to generate the 10 seed videos used in our research. Recommended videos were gathered using custom-made crawlers over "depths" of recommendations. The seed videos generated the 1st video depth, after which subsequent depths served as parent videos to generate the next sets of recommended videos. This process continued until recommended videos for 5 depths were generated. To prevent personalization in recommendations, we did not log into the account used for video collection. Also, a new browser instance was started and cookies from each previous recommendation depth were cleared to enable a fresh search of videos for the next depth crawl. This approach allowed us to generate a total of 38,970 videos across 5 depths. To focus our study, we filtered out duplicates and non-English videos which reduced our dataset to 14,332 videos, after which videos were categorized by depth. Video text data such as titles, descriptions and transcripts were used for this research.

The collection of video transcripts was divided into four sub-tasks.**Task 1:** Video transcripts were fetched using YouTube Transcript API [24]. 14,332 video ids were fed to the API and 12,611 transcripts were gathered. **Task 2:** We found that transcripts were disabled for 1,721 videos. For such videos, we used the OpenAI Whisper model [20] to extract the video transcripts. This led to an additional 1567 transcripts of which 154 videos were unavailable as they were identified as 'live shows', 'removed' or returned null in our script. **Task 3:** We identified and translated non-English transcripts to English transcripts using Google Translate API and removed transcripts which had less than 80% English content. **Task 4:** Lastly, the results were combined together and processed for analysis.

IV. Methodology

In this section, we discuss the techniques used in our study.

A. Emotion and Popularity Assessment Methodology

For this study, we analyzed emotions embedded in video text data (title, description and transcript) across 6 emotions: joy, anger, sadness, fear, surprise, love. We use emotion drift to identify emotion bias across depths of recommendations. The resulting emotion diversity in content were illustrated on a line graph with each depth representing a traversed hop of recommended videos. A fine-tuned version of transfer learning [10], T5-base-fine-tuned-emotion, was utilized for Natural Language Processing (NLP) tasks to ensure accuracy of results. To further understand the emotion drift pattern in recommended videos, we analyzed user engagement using engagement metrics of all videos such as likes and views. With the engagement metrics, we studied the change in metrics across depths, to determine if the user interaction supports the emotion drift pattern across recommendation depths.

B. Topic, Network and Content Analysis Methodology

Although previous research methodologies have concatenated video text information (video titles, video description, and video transcript) for video content analysis [18], this research analyzed these three components separately as well as in combination. By analyzing these components separately, we hoped to identify a variability in content concentration at varying levels of video text detail. The goal of topic drift detection is to investigate if recommendations stay on the topic of the Uyghur crisis as we move through recommended videos and by how much content diverges if drift is detected. To measure topic drift across depths, we computed topic similarity using Hellinger distance [11], [12], [29] and Jensen-Shannon divergence [26], [28]. Hellinger and Jensen-Shannon divergence are distance metrics used in estimating document similarity. Hellinger divergence is represented as the symmetric midpoint of Kullback–Leibler divergence [25], [30] while Jensen-Shannon divergence is a finite, smoothed version of Kullback–Leibler divergence [27]. These distance metrics calculate similarity within the range of 0 to 1, where values closer to 0 indicate a smaller distance and, therefore, larger similarity. We computed a final topic similarity score using the average of both scores across depths. Next, we analyzed the video recommendation network. Recommendation graphs for each depth consisting of video ids as nodes and recommendations as edges were generated and examined. The distribution of eigenvector centrality scores, which measure the influence a node has on a network of videos across depths was computed to determine if a sub-cluster of videos were highly influential (more recommended) compared to other videos. We then analyzed our data to determine the topic communities of videos in each recommendation depth. For this approach, we used the BertTopic model [14], a model which uses transformers and class-based term frequency-inverse document frequency (c-TF-IDF) to create dense clusters and produce interpretable topics, while maintaining important words in the topic description. By programmatically assigning each video to its respective topic community across depths. We were able to detect how influential videos evolved across recommendations.

V. Results

A. Topic Drift Analysis

As earlier discussed, the goal of topic drift detection is to determine if recommended videos stay on the topic of the China-Uyghur crisis or drift as users move through recommended videos. Distance metrics are often measured between 0 and 1, where scores closer to 0 depict high similarity (contents are similar) and scores closer to 1 depict low similarity (contents are different). For this research, drift is observed if there is an increase in the distance between depths resulting in decreased content similarity. This is seen as a rising trend in the distance metric line graph. Using video titles, video descriptions, video transcripts, and a concatenation of all text information to analyze topic drift, we compared the similarities to answer two questions;

- **Are our seed videos different from videos across depths?**

This question was answered by analyzing the similarity between the seed videos and each depth of recommendation. The goal was to measure the video similarity between the seed and recommended videos in each depth of recommendation.

- **Do recommended videos become increasingly similar or different from each other?**

This question was answered by analyzing the similarity between adjacent depths of recommendation. The goal was to measure the video similarity across depths of recommended videos.

1) Similarities between the seed videos and subsequent depths of recommendation: In Fig. 1, we observe that the seed videos are significantly different from depth 1 videos. Once we approach depths 2 – 3, the videos increase in similarity to the seed videos compared to videos in depth 1, but remain significantly different in general. This result shows that, depth 1 recommendations were highly unrelated to the China-Uyghur crisis, but, as the users moves through depths 2 to 5, the videos become somewhat similar to our seed videos, but not to a relevant degree.

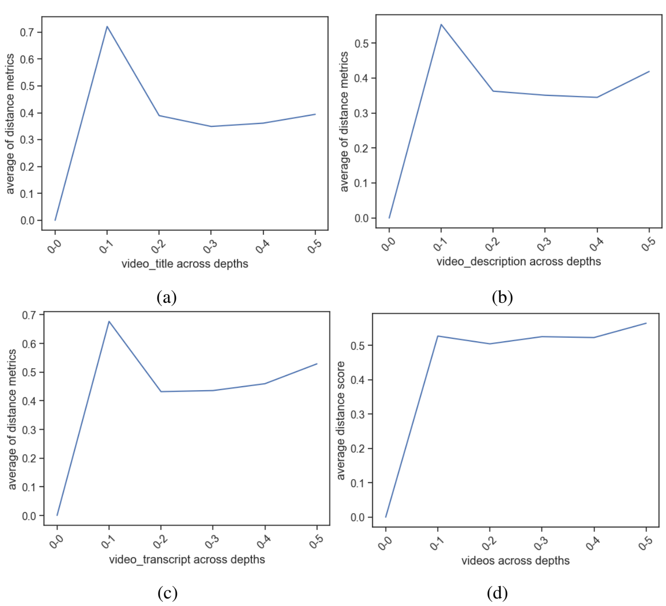

(a) (b) (c) (d)

Fig. 1: Line graph showing how recommended videos become increasingly different from seed videos using (a) video titles (b) video descriptions (c) video transcript (d) all text information

2) Similarities between adjacent depths of recommendation: This question investigates how similar each depth of recommended videos is compared to its previous depth. From this analysis, we observe that as we move through recommended videos, each depth of videos becomes more similar to its previous depth, reaching maximum similarity between depths 3 and 4. Both results suggest that, although each depth of recommended videos becomes more different from our seed videos, each depth of videos also becomes more similar to its previous depth. With this pattern, the difference in recommended videos is not immediately noticed and the user is gently re-introduced to content unrelated to the seed videos.

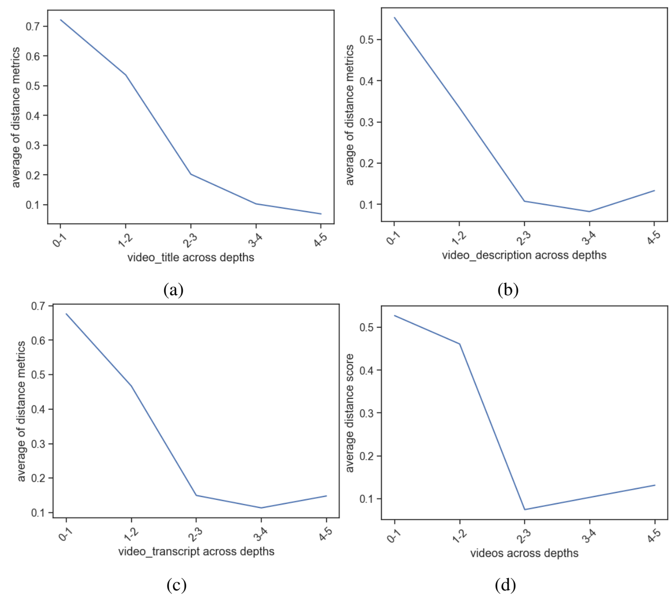

(a) (b) (c) (d)

Fig. 2: Line graph showing how recommended videos become more similar across depths using (a) video titles (b) video descriptions (c) video transcript (d) all text information

B. Network and Content Analysis

Next, network analysis was performed on each depth of recommended videos. For each depth, each video is ranked using its eigenvector centrality measure, to determine its influence in the network. For a given graph G:=(V,E) with —V— vertices, let A = (avt) be the adjacency matrix, i.e., avt = 1 if vertex v is linked to vertex t, and avt = 0 otherwise. The relative centrality score, Xv of vertex v can be defined as:

$$X_y = \frac{1}{\lambda}\Sigma_{t\epsilon M(v)}X_t = \frac{1}{\lambda}\Sigma_{t\epsilon v}a_{vt}X_t \tag{1}$$

where M(v) is the set of neighbors of v and is a constant. With a small rearrangement, this can be rewritten in vector notation as the eigenvector equation.

$$Ax = \lambda x \tag{2}$$

To find the most influential videos, we isolated and analyzed the top 10 videos with the highest eigenvector centrality score per depth. The mean eigenvector centrality score for the top 10 videos per depth was found and videos which had an eigenvector centrality score above the resulting mean were identified and categorized as ‘above-average’ / highly influential videos. Our results suggest that these ‘above-average’ influential videos act as attractors by driving the recommendations of videos and directing how the conversation evolves across depths. We also see that the top 10 videos in each depth fluctuate in the count of ‘above-average’ influential videos in each depth, as seen in Table I. As we move through the depths, we observe a steady increase in the number of ‘above-average’ influential videos until depth 3. Once we arrive at depth 3, the count of above-average influential videos began to steadily decrease. Also, content of these influential videos seem to drift from our seed videos as we approach depth 5. To visualize the content divergence of above-average videos from our seed videos after depth 3, we performed topic modelling on the seed videos and the whole dataset to generate the latent topics present in the recommendations and assign each video a topic community number.

TABLE I: TOPIC COMMUNITIES OF HIGHLY-INFLUENTIAL VIDEOS PER DEPTH

Videos	Count of highly-influential videos	Topic communities
Seed	N/A	-1
Depth 1	1	1
Depth 2	4	-1, -1, -1, -1
Depth 3	4	1, 9, -1, -1
Depth 4	2	1, 9
Depth 5	3	1, 9, 13

Topic modelling was done using BERTopic, to identify the topic communities present in our seed videos and the topic communities of highly influential videos in each recommendation depth. By doing this, we were able to visualize the topical content of our ‘above-average’ influential videos and track the movement of content topically related to our seed videos as we moved across depths. We observed that all of our seed videos belonged to one topic community, -1, while the highly influential videos across depths belonged to a mix of topic communities. From Table I, we see that the highly influential video at depth 1 is introduced into the algorithm, and steers depth 1 away from the content on Uyghur crisis. Conversely, as we move to depth 2 the highly influential videos are fully turned back to topics related to our seed videos. At depth 3, the highly influential videos contain an equal mix of videos related and unrelated to our seed videos but once we arrive depth 4, our seed video content is filtered out once more from the list of highly influential videos. This result shows that, as we progress through the recommendations, videos related to our seed videos are filtered out from the recommendations in a pendulum-like motion. From Table I, we observe that the algorithm seems to swing back and forth from content related to the Uyghur crisis, reducing its influence with each motion until it is finally filtered out of the recommended videos. We are also able to see that content in depth 5 is topically unrelated to our seed videos as seen in the difference in topic communities from our seed videos in Table I.

C. Emotion and Popularity Analysis

1) Emotion Analysis: To study the pattern of emotion drift across depths, we considered video text data at 4 different levels; video titles, video description, video transcript and a combination of all texts. By doing this, we were able to apply emotion assessment and visualize emotion drift at different levels of video details, as seen in Fig. 3(a), 3(b), 3(c) and 3(d). The results show that the most dominant emotion in our seed videos was anger for all levels of video detail, as illustrated in the figures. As we traverse the recommendation depths, we see the positive emotion (joy) emerge for each depth in all emotion graphs and the negative emotions (anger, fear, and sadness) decrease significantly.

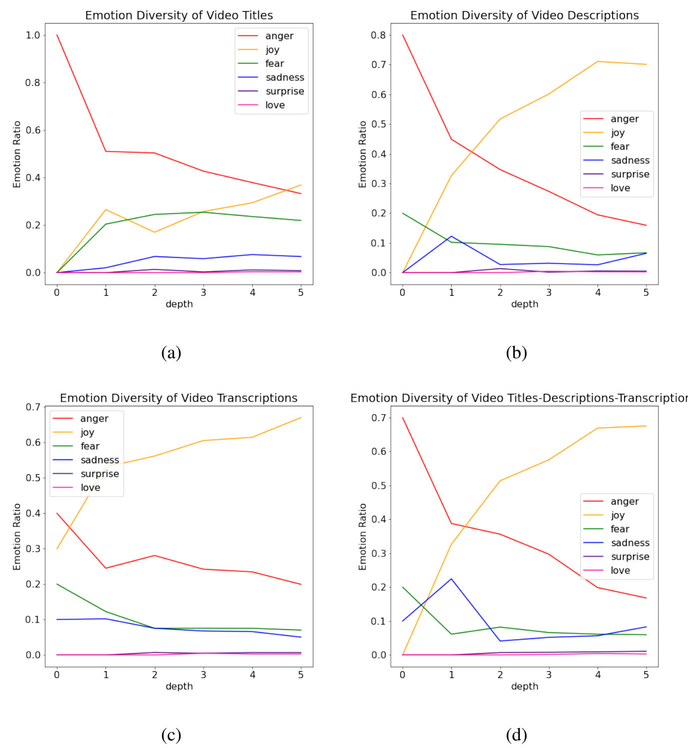

(a) (b) (c) (d)

Fig. 3: Emotion assessment for video text data (a) titles only (b) descriptions only (c) transcripts only (d) all text information

2) Popularity Analysis: By analyzing the emotions of the video across depths using video text data, we discovered that there was a significant decrease in negative emotion (anger) and a significant increase in positive emotion (joy). To investigate the significance of this emotion drift pattern, we analyzed user interaction with the videos using engagement metrics across depths. This analysis was to determine if more popular videos were recommended across depths. For this experiment, a popular video is described as a video which has significantly high views and high positive engagement in the form of likes. As a result, the engagement metrics we considered were the views and likes of each video. On inspecting our seed videos, we found they all had a very high view count but a significantly low like count, suggesting that although our seed videos were widely watched, they did not elicit positive interaction from the audience. This is to be expected as the China-Uyghur crisis has been monitored internationally with the discourse being widely criticized. From the video like box-plot in Fig. 4(a), we see that, as we move through recommendation depths, the median likes of recommended videos are significantly higher compared to the seed videos and increase linearly until we hit depth 3, after which, there is an exponential increase in video likes by depth 4 and depth 5. Secondly, our video views box-plot in Fig. 4(b) shows that the views of recommended videos are higher compared to the seed videos but unlike video likes, we see a steady growth in view count across depths of recommended videos. The result of our popularity analysis shows that more popular videos are present in recommended videos, which further explains the high occurrence of positive emotions in higher depths of recommendations.

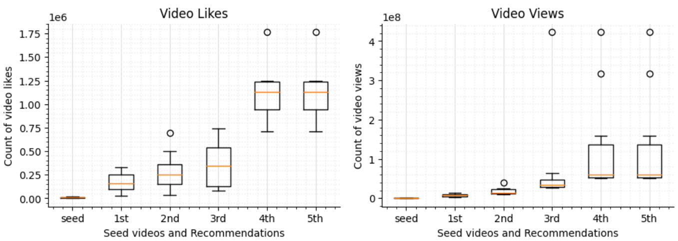

(a) (b)

Fig. 4: The box-plot show the increasing median count of (a) video likes and (b) video views from seed videos to recommended videos.

VI. Discussion

RQ1: How do we identify bias on content related to vulnerable communities?

In examining the results from our emotion and popularity analysis, we observed that the anger emotion significantly decreases across depths, while there is a proportional increase in the joy emotion in videos after each recursive depth of recommendation. In addition, we see that the engagement metrics (views and likes) increase significantly as we move to higher depths of recommendation, suggesting increased user engagement with recommended videos. In summary, the algorithm seems to recommend more popular videos with positive emotions (joy) in an attempt to keep users engaged for longer periods of time. This pattern demonstrates recommender bias which steers users away from unpopular videos with negative emotions. This trend poses the risk of occluding content related to the China-Uyghur crisis.

RQ2: How do videos related to vulnerable communities change across recommendation depth?

Our topic drift analysis shows that as users watch recommended videos, the videos become increasingly different from our seed videos across recommendations. We also found that each depth of recommended videos became increasingly similar to its immediate previous depth suggesting that videos across recommendations are similar in content. These drift patterns show that the algorithm gently drifts from our seed videos by recommending videos that are increasingly different from our seed videos but similar to adjacent depths of recommendations until recommended videos significantly drift from content related to the China-Uyghur crisis at depth 5.

RQ3: What kind of videos drive recommendations in the context of this study?

Through our network analysis, we observe that each depth has a set of highly influential videos which act as attractors to drive video recommendations. The gradual shift in topics we observe from seed videos to depth 5 in Fig. 1 seems to be due to a *pendulum-like* motion of the algorithm. From Table I, our results show that depths 3, 4 and 5 show a back-and-forth swing of the algorithm. There is an alternate filtering and re-introduction of content related to our seed videos across depths, maintaining a steady plateau in similarity of depths 3 - 5 to our seed videos until the China-Uyghur crisis topics are filtered out of the recommendations.

VII. Conclusion and Future Works

For this research, we employed the use of drift analysis to identify bias across recommended videos. Our results showed that YouTube's recommendation system tends to lessen negative emotions such as anger and amplify positive emotions such as joy across recommended videos on the platform. We also see that highly influential videos at each depth act as attractors to gently draw recommendations away from content related to our seed videos in a pendulum-like motion. In future research, we plan to expand this research into exploring an alternate narrative which elicits a different emotion (e.g joy) and comparing the findings with those of our current research. We are also interested in developing a framework which serves to methodologically compare content across various discourse and exploring the effects of the YouTube algorithms on such datasets.

Acknowledgment

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] S. N. Shivhare and S. Khethawat, "Emotion Detection from Text." arXiv, May 22, 2012. doi: 10.48550/arXiv.1205.4944.

[2] J. M. Garcia-Garcia, V. M. R. Penichet, and M. D. Lozano, “Emotion detection: a technology review,” in Proceedings of the XVIII International Conference on Human Computer Interaction, New York, NY, USA, Sep. 2017, pp. 1–8. doi: 10.1145/3123818.3123852.
[3] J. E. Solsman, “Ever get caught in an unexpected hourlong YouTube binge? Thank YouTube AI for that,” CNET, Jan. 10, 2018. https://www.cnet.com/tech/services-and-software/youtube-ces-2018-neal-mohan/ [Accessed Jan. 28, 2023].
[4] D. M. Blei, A. Y. Ng, and M. I. Jordan, “Latent dirichlet allocation,” J. Mach. Learn. Res., vol. 3, no. null, pp. 993–1022, Mar. 2003.
[5] Q. Liu, H. Huang, and C. Feng, “Micro-blog Post Topic Drift Detection Based on LDA Model,” in Behavior and Social Computing, Cham, 2013, pp. 106–118
[6] F. A. Acheampong, C. Wenyu, and H. Nunoo-Mensah, “Text-based emotion detection: Advances, challenges, and opportunities,” Engineering Reports, vol. 2, no. 7, p. e12189, 2020, doi: 10.1002/eng2.12189.
[7] N. O’Hare et al., “Topic-dependent sentiment analysis of financial blogs,” in Proceedings of the 1st international CIKM workshop on Topic-sentiment analysis for mass opinion, New York, NY, USA, Nov. 2009, pp. 9–16. doi: 10.1145/1651461.1651464.
[8] M. Suhasini and S. Badugu, “Two Step Approach for Emotion Detection on Twitter Data,” International Journal of Computer Applications, vol. 179, pp. 12–19, Jun. 2018, doi: 10.5120/ijca2018917350.
[9] A. Murakami, P. Thompson, S. Hunston, and D. Vajn, “‘What is this corpus about?’: using topic modelling to explore a specialised corpus,” Corpora, vol. 12, no. 2, pp. 243–277, Aug. 2017, doi: 10.3366/cor.2017.0118.
[10] C. Raffel et al., “Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer,” Journal of Machine Learning Research, vol. 21, no. 140, pp. 1–67, 2020.
[11] A. Shemyakin, “Hellinger Distance and Non-informative Priors,” Bayesian Analysis, vol. 9, no. 4, pp. 923–938, Dec. 2014, doi: 10.1214/14-BA881.
[12] G.-H. Fu, Y.-J. Wu, M.-J. Zong, and J. Pan, “Hellinger distance-based stable sparse feature selection for high-dimensional class-imbalanced data,” BMC Bioinformatics, vol. 21, no. 1, p. 121, Mar. 2020, doi: 10.1186/s12859-020-3411-3.
[13] O. Papakyriakopoulos, J. C. M. Serrano, and S. Hegelich, “Political communication on social media: A tale of hyperactive users and bias in recommender systems,” Online Social Networks and Media, vol. 15, p. 100058, Jan. 2020, doi: 10.1016/j.osnem.2019.100058.
[14] M. Grootendorst, “BERTopic,” Mar. 11, 2022. https://maartengr.github.io/BERTopic/ [Accessed Jan. 09, 2023].
[15] “China’s Repression of Uyghurs in Xinjiang — Council on Foreign Relations.” https://www.cfr.org/backgrounder/china-xinjiang-uyghurs-muslims-repression-genocide-human-rights [Accessed Jan. 09, 2023].
[16] H. Heuer, H. Hoch, A. Breiter, and Y. Theocharis, “Auditing the Biases Enacted by YouTube for Political Topics in Germany,” in Proceedings of Mensch und Computer 2021, New York, NY, USA, Sep. 2021, pp. 456–468. doi: 10.1145/3473856.3473864.
[17] K. Topal, M. Koyuturk, and G. Ozsoyoglu, “Emotion -and area-driven topic shift analysis in social media discussions,” in 2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), Aug. 2016, pp. 510–518. doi: 10.1109/ASONAM.2016.7752283.
[18] M. Faddoul, G. Chaslot, and H. Farid, “A Longitudinal Analysis of YouTube’s Promotion of Conspiracy Videos.” arXiv, Mar. 06, 2020. doi: 10.48550/arXiv.2003.03318.
[19] L. P. Goldsmith et al., “The use of social media platforms by migrant and ethnic minority populations during the COVID-19 pandemic: a systematic review.” medRxiv, p. 2022.02.07.22270579, Feb. 07, 2022. doi: 10.1101/2022.02.07.22270579.
[20] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust Speech Recognition via Large-Scale Weak Supervision.” arXiv, Dec. 06, 2022. doi: 10.48550/arXiv.2212.04356.
[21] B. Auxier, “Social media continue to be important political outlets for Black Americans,” Pew Research Center. https://www.pewresearch.org/fact-tank/2020/12/11/social-media-continue-to-be-important-political-outlets-for-black-americans/ [Accessed Jan. 09, 2023].
[22] C. Silverman, “This Analysis Shows How Viral Fake Election News Stories Outperformed Real News On Facebook,” BuzzFeed News. https://www.buzzfeednews.com/article/craigsilverman/viral-fake-election-news-outperformed-real-news-on-facebook [Accessed Jan. 09, 2023].
[23] B. Kitchens, S. L. Johnson, and P. Gray, “Understanding Echo Chambers and Filter Bubbles: The Impact of Social Media on Diversification and Partisan Shifts in News Consumption,” Aug. 26, 2020. https://misq.umn.edu/understanding-echo-chambers-and-filter-bubbles-the-impact-of-social-media-on-diversification-and-partisan-shifts-in-news-consumption.html [Accessed Jan. 09, 2023].
[24] J. Depoix, “YouTube Transcript/Subtitle API (including automatically generated subtitles and subtitle translations).” Jan. 08, 2023. [Accessed: Jan. 09, 2023]. [Online]. Available: https://github.com/jdepoix/youtube-transcript-api
[25] H. Sengar, H. Wang, D. Wijesekera, and S. Jajodia, “Detecting VoIP Floods Using the Hellinger Distance,” IEEE Transactions on Parallel and Distributed Systems, vol. 19, no. 6, pp. 794–805, 2008, doi: 10.1109/TPDS.2007.70786.
[26] M. Jamaati and A. Mehri, “Text mining by Tsallis entropy,” Physica A: Statistical Mechanics and its Applications, vol. 490, pp. 1368–1376, Jan. 2018, doi: 10.1016/j.physa.2017.09.020.
[27] A. Mehri, M. Jamaati, and H. Mehri, “Word ranking in a single document by Jensen–Shannon divergence,” Physics Letters A, vol. 379, no. 28, pp. 1627–1632, Aug. 2015, doi: 10.1016/j.physleta.2015.04.030.
[28] D. M. Endres and J. E. Schindelin, “A new metric for probability distributions,” IEEE Transactions on Information Theory, vol. 49, no. 7, pp. 1858–1860, 2003, doi: 10.1109/TIT.2003.813506.
[29] S. Sohangir and D. Wang, “Improved sqrt-cosine similarity measurement,” Journal of Big Data, vol. 4, no. 1, p. 25, Jul. 2017, doi: 10.1186/s40537-017-0083-6.
[30] S. Zhu, L. Liu, and Y. Wang, “Information retrieval using Hellinger distance and sqrt-cos similarity,” in 2012 7th International Conference on Computer Science and Education (ICCSE), 2012, pp. 925–929. doi: 10.1109/ICCSE.2012.6295217.