Unveiling Content Traps: A Network Resilience and Topic-Based Study of YouTube's Algorithmic Content Curation

Md Monoarul Islam Bhuiyan*, Nitin Agarwal*†
*COSMOS Research Center, University of Arkansas at Little Rock, Arkansas, USA
†International Computer Science Institute, University of California, Berkeley, California, USA
e-mail: {mbhuiyan, nxagarwal}@ualr.edu

***Abstract*—YouTube's recommendation algorithm drives 70% of total video views, playing a pivotal role in shaping user engagement and content consumption. However, this algorithm often contributes to the creation of content traps, or attractor content, where users are nudged towards certain videos. Repeated nudges to users in turn create filter bubbles and echo chambers. This study aims to identify and analyze such content traps within YouTube's recommendation network using Focal Structure Analysis (FSA), a social network analysis approach designed to detect key groups of nodes in a network that together act as powerful attractors due to their position in the network. Such focal structures in a video recommendation network act as strong content traps. By constructing a network based on YouTube's recommendation hops, we identify content traps using the FSA approach. We then confirm the presence of dominant themes and high topical consistency by leveraging topic modeling and information divergence metrics, suggesting less topical diversity in content traps. Engagement analysis of content trap videos with high topic uniformity reveals user interactions reinforce these traps. Our findings provide insight into content trap dynamics that could inform strategies for content consumers to break out of them and help platform owners to develop debiasing techniques in their content curation algorithms. This research highlights the critical role of recommendation algorithms in shaping content exposure and calls for strategic interventions to foster greater content diversity and mitigate the effects of content traps.**

Keywords-Content Traps; Network Resiliency; YouTube Recommendation Network; Social Network Analysis; BERTopic

I. Introduction

With the rise of social media platforms, several activities have gained wide popularity, including content creation and sharing, news consumption, community engagement, societal influence, and many others. The rapid adoption of these has become possible due to massive user engagement with video content that can often lead to the entrapment of content [1]. YouTube is the second-most popular social media platform and the number one video-sharing platform globally, available in over 100 countries and 80 languages [2]. Moreover, YouTube's recommendation algorithm drives 70% [3] of its watch time, making it one of the most influential aspects of the platform. Since this algorithm helps users discover content aligned with their viewing habits, it often also creates content traps or attractor content, where users are repeatedly exposed to similar content. The emergence of content traps is particularly concerning for sensitive or controversial topics, as it can influence public opinion, deepen polarization, and propagate misinformation. Given YouTube's critical role in shaping digital discourse, understanding the mechanisms driving content traps is essential. This phenomenon highlights concerns about the recommendation algorithm's influence on user behavior and its role in reinforcing content loops.

With an emphasis on how YouTube's algorithm directs users toward repeating and thematically limited content loops, we examine in this study the phenomena of content traps inside the platform's recommendation network. We employ Focal Structure Analysis (FSA) [4], a social network analysis (SNA) approach, to detect powerful set of nodes, known as focal structures, that play a significant role in shaping user engagement and reinforcing content traps.

Our analysis begins by constructing a recommendation network based on video suggestion patterns and then identifying focal structures. After that, we assess the impact of these structures by examining how their removal from the recommendation network affects its overall connectivity and resilience.

Through topic modeling and divergence metrics, we explore the topical consistency within these focal structures. Our study identifies content traps characterized by dominant, repetitive themes. Moreover, we examine user engagement trends to determine how interactions contribute to the persistence of these traps. Our study highlights how crucial it is to address algorithmic design in order to encourage content diversity and lessen the effects of content loops, which will ultimately help create a more equitable online interaction. This research highlights the broader implications of recommendation systems on user behavior and content consumption in online platforms.

Although algorithmic biases are becoming more widely recognized [5], to the best of our knowledge, there is currently no systematic method for identifying and assessing these content traps within YouTube's dynamic recommendation system. Our study aims to understand the presence of content traps within YouTube's recommendation network by leveraging the focal structure analysis approach. With this, we seek to answer the following research questions:

- **RQ1**. How can content traps or attractor contents be detected within focal structures from YouTube's recommendation network?
- **RQ2.** How can topical consistency within focal structures of YouTube's recommendation network be assessed to identify the presence of content traps or attractor contents?
- **RQ3.** Can engagement statistics be used to validate the findings of the content trap discovery model?

We examine how focal structures—defined as key sets of attractor nodes within YouTube's recommendation network—contribute to the emergence and persistence of content traps. By analyzing how these structures shape user experiences and exposure pathways, we reveal how recommendation algorithms can narrow informational diversity and reinforce homogeneous content loops. This study advances the understanding of algorithmic content traps by integrating structural, semantic, and behavioral perspectives. We show that Focal Structure Analysis systematically identifies cohesive subgraphs that act as attractor contents, that these structures display strong thematic uniformity as measured through BERTopic and divergence metrics, and that user engagement disproportionately amplifies the most uniform topics, empirically validating trap-like behavior. Together, these insights contribute to the broader discourse on algorithmic content control and offer a generalizable analytical approach for diagnosing and mitigating entrapment dynamics within large-scale recommendation platforms.

The rest of this paper is organized as follows: The Literature Review section reviews existing studies for identifying influential sets of nodes from a social network along with the identification and implications of content traps in social media. While the Methodology section outlines the approaches used for the overall study, the Experimental Results and Analysis section presents the findings of our research. Lastly, the final Conclusion and Future Work section summarizes the study and suggests directions for future research.

II. Literature Review

This section is divided into two parts. We first discuss the relevant literature related to identifying influential sets of nodes within a social network. After that, we delve deeper into the related literature around content traps and topical uniformity [6], exploring their formation and impact on social networks as well as the implications for user experience and societal impact from these attractor contents in social media.

A. Identifying Sets of Key Social Network Entities

Identifying key nodes or sets of nodes that are best connected and most influential in a social network is crucial for extracting actionable knowledge. Consequently, various methods have been proposed to identify these key nodes. While HITS determines hubs and authorities [7], PageRank assigns a numerical weight for each node in the network [8]. Both of these approaches can be used to identify influential nodes from social networks. In contrast, identifying the communities and clusters from a social network perspective has also been widely studied. Generally, in a community, similar nodes are more clustered together than nodes that don't share commonalities [9]. Previous researchers have also worked on a more sophisticated approach where their focus shifted from identifying the influential nodes or communities to detecting smaller key sets of players who maximized information diffusion. Authors in [10] devised a methodology where they identified focal patterns leveraging the Louvain method, which gave them more relevant information about the network than obtained from just the influential nodes [11]. When applying this method to large biological networks, they found more prominent, smaller, and relevant structures in protein-protein interaction networks [12]. Since this method couldn't extract structures with lower connection density, researchers extended their approach by combining highly connected candidate focal structures based on similarity values. This allowed the identification of both cliquish and small sparse yet connected structures [13]. An advanced version of this approach was proposed by [4], where they combined user-level centrality and group-level modularity methods to create a bi-level maximization network model that overcame the shortcomings of the previously described focal structures analysis methods.

While identifying these nodes is crucial for network analysis since they greatly influence the network's architecture, their impact assessment on the overall network is also imperative. As a consequence several of such assessment metrics have been developed [14] to quantify the network's resilience to disruptions, providing insights into its stability and connectivity.

B. Content Homogeneity in Social Media

Content traps occur when recommendation systems such as those used by social media platforms limit users' exposure to diversified content, which in turn can result in the reinforcement of existing beliefs and preferences. Pariser et al. (2011) first introduced the concept of filter bubbles, highlighting how algorithms can create information ecosystems that reinforce users' existing views and interests. More recent research by Bechmann et al. (2018) examined filter bubbles in Facebook news feeds by analyzing information similarity in link sources and content semantics. The role of recommender systems in increasing content homogeneity, particularly during significant political events such as the 2018 Brazilian presidential election, has also been studied. While the study of content bubbles has been conducted, solutions to mitigate these bubbles have also been researched. Diversification algorithms have been explored as a potential solution to mitigate these election polarization effects [17]. Studies on link prediction in social networks show that fairness-aware algorithms, which utilize network modularity measures, can decrease the segregation caused by homophily and mitigate filter bubbles [18]. On a behavioral level, awareness of filter bubbles influences how users respond to them, though personality traits have been found to play a less significant role in motivating corrective actions, particularly on platforms like Facebook [19]. Furthermore, an integrated tool synthesizing previous research has been proposed to provide users with a more comprehensive way to avoid filter bubbles across social networks, emphasizing the need for holistic approaches to diversify user experience and counteract content loops [20]. Together, these studies emphasize the growing importance of addressing filter bubbles through technological, algorithmic, and user-behavior-focused interventions.

Despite substantial attention to filter bubbles and algorithmic bias, there remains limited empirical work that systematically identifies and quantifies content traps within YouTube's recommendation network. Existing studies often address influential nodes, content homogeneity, or user-level effects in isolation, but they do not provide an integrated framework that connects structural attractors, semantic uniformity, and engagement-driven reinforcement. This fragmentation leaves a gap in understanding how recommendation algorithms produce and sustain narrow content pathways. In this paper, we address this gap by identifying focal structures within YouTube's recommendation graph, assessing their structural influence through network fragmentation analysis, and evaluating their thematic consistency using topic modeling and divergence metrics. By linking these structural and semantic findings with engagement patterns, our study offers a unified approach for understanding how content traps emerge and persist within algorithmically curated environments.

III. Methodology

This section outlines our systematic approach to analyzing the YouTube recommendation network, detecting focal structures, and evaluating the presence of content traps. We start by summarizing our approach in collecting data, dataset background, and building YouTube recommendation networks. After that, we present the network resiliency approach taken to rank key focal structures. In addition to that, we lay out the foundation for the analysis of the topics using the BERTopic model. Lastly, this section explores several metrics to investigate the topical consistency across different topics within the focal structures.

A. Data Collection

The data collection process in this study was designed to systematically capture YouTube's algorithmic behavior through its 'watch-next' recommendations. In this study. we analyzed two contexts, the China–Uyghur conflict and the Cheng Ho propaganda on YouTube. Below we provide background details for these two contexts and the motivation for studying them.

1) Cheng Ho Propaganda: In recent years, the Chinese Communist Party (CCP) has adapted the story of the 15th-century admiral Zheng He, also known as Cheng Ho, to support its current political messages. Once known for his peaceful sea voyages, Zheng He is now depicted as a symbol of religious tolerance and diplomacy [21]. This shift aligns with China's efforts to address criticism of its treatment of Uyghur Muslims and to promote its Maritime Silk Road initiative. By rebranding this historical figure, the CCP aims to strengthen its soft power, particularly in Southeast Asia.

2) China Uyghur Conflict: The conflict in Xinjiang focuses on the struggles faced by the Uyghur Muslim minority in China, which includes issues like cultural suppression, ethnic tensions, and government policies[22]. Researchers have studied this situation through various lenses, such as identity politics, language policies, the dynamics between majority and minority groups, and the desire for self-governance [23]. The international response between 2018 and 2022 has drawn more attention to human rights issues, highlighting the severity of the conflict.

We selected the China–Uyghur conflict and Cheng Ho propaganda datasets for their geopolitical and ideological relevance in examining algorithmic content amplifications and recommendation dynamics within the recommendation network.

3) Keyword Generation and Crawling: We initiated the process by conducting workshops with subject matter experts to compile a targeted list of keywords related to the China–Uyghur conflict and the Cheng Ho propaganda. These keywords served as search queries on YouTube's search engine, generating an initial set of seed videos. It's important to note that the unbalanced count of keywords as shown in Table I and Table II between different datasets does not compromise the validity of our study.

Table I: Keywords related to China-Uyghur conflict

Keywords
Penindasan/oppression + Uighur/Uyghur, Kejam/cruel + Uighur/Uyghur, Saudara muslim/muslim brother + Uighur/Uyghur, Kalifah/capliph + Uighur/Uyghur, Khilafah/caliphate + Uighur/Uyghur, "China is Terrorist", "Stop Genocide", "Save Muslim Uyghur", "Get Out China", "I Love Muslim Uyghur", "peduli Uyghur" / "Care Uyghur", "bebaskan muslim uyghur dari penindasan china" / "free uyghur muslims from china's oppression", "do'a kan saudaramu" / Pray for Muslim Uyghur, Hizbul Tahrir / HTI + Uighur/Uyghur, Front Pembela Islam/FPI/Islamic Defenders Front + Uighur/Uyghur, Nahdlatul Ulama + Uighur/Uyghur, Muhammadiyah + Uighur/Uyghur, Hebibulla Tohti + Indonesia, Mohammed Salih Hajim + Indonesia, Yusuf Martak + Uighur/Uyghur, Slamet Ma'arif + Uighur/Uyghur, Xiao Qian + Uighur/Uyghur, Pendidikan/education + Uighur/Uyghur

Table II: Keywords related to Cheng Ho propaganda

Keywords
Cheng Ho, Zheng He, Sam Po Kong, Sam Poo Kong, Daerah Otonom Uighur Singkiang, Singkiang, Hatta + 1957, Novi Basuki, Sam Po Bo, Cheng Ho / "Zheng He" + laksamana + damai, Sam Po Kong + Islam + Indonesia, 1421 Saat China Menemukan Dunia + "Gavin Menzies", Gavin Menzies, Cheng Ho / "Zheng He" "Columbus"

Using custom-built crawlers, video recommendations were extracted recursively across multiple depths. Starting with the seed videos, the first depth of recommendations was captured, and each recommended video was subsequently treated as a parent node for further investigation [24]. Datasets with distinct videos were produced due to this iterative procedure, which proceeded until four recommendation depths were reached. This created the recommendation network with 5 levels. The decision to limit the collection of data for 5 levels was a strategic choice based on essential computational constraints. As we explored deeper recommendation levels, the data volume increased exponentially. Restricting the analysis to a maximum of 5 levels for a more manageable evaluation of the data without compromising the integrity of the findings. A similar strategy was adopted by the authors in [25]. We then gathered titles, descriptions, and engagement metrics such as likes, comments, and views for each video except for the transcript using YouTube Data API v3 [26]. However, for transcript generation, the method from [27] was used.

B. Network Construction

The China–Uyghur conflict and Cheng Ho propaganda datasets consisted of 9,748 videos and 14,307 interconnections and 8,489 videos and 13,384 interconnections, respectively, that represented YouTube's recommendation paths across several hops, and they were used to start the network construction process. The initial networks had an average clustering coefficient of 0.067 and 0.062, respectively, suggesting comparatively sparse and loosely connected network structures.

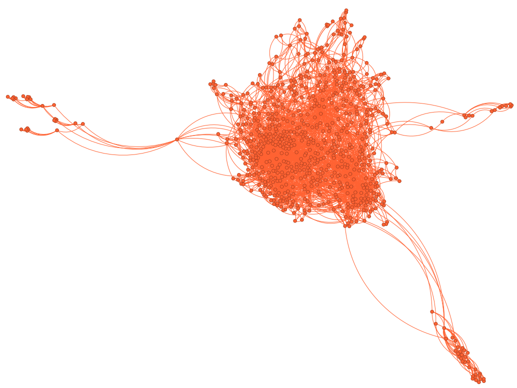

Figure 1: This graph illustrates the Cheng Ho propaganda recommendation network on YouTube after pruning, showing a denser, more interconnected structure with an average clustering coefficient of 0.362.

Figure 2: This graph illustrates the China–Uyghur conflict recommendation network on YouTube after pruning, showing a denser, more interconnected structure with an average clustering coefficient of 0.38.

To make our study more efficient and less computationally constrained, we then refined the recommendation networks using a filtering criterion to concentrate on significant connections by retaining only videos that had a degree of 3 or more. This simple yet effective graph-pruning approach reduced network sparsity, enhanced structural coherence, and ensured that the analysis concentrated on nodes that were more influential, aligning with prior research that had demonstrated the effectiveness of graph pruning in improving computational feasibility and analytical precision [28], [29]. Thus, the filtering process reduced the networks, resulting in a denser graph with an average clustering coefficient of 0.38 and 0.362, respectively, for the China–Uyghur conflict and Cheng Ho propaganda recommendation networks, as depicted in Figures 1 and 2. The increase in the clustering coefficient indicates a more tightly-knit YouTube recommendation network, ensuring that our analysis captured well-connected focal structures rather than weakly linked, peripheral nodes. Due to this, we identified areas more likely to contribute to content traps by focusing on these denser network connections. This allowed us to assess the impact of focal structures on the topical consistency of recommended content and their influence on user engagement.

C. Focal Structure

Focal structures (FS) are key sets of individuals in a social network who may be responsible for coordinating events, protests, or leading citizen engagement efforts. In the context of content traps, focal structures can be defined as a set of attractor content(s), e.g., a set of videos. By manipulating content recommendations, these structures limit exposure to diverse perspectives, reinforcing particular viewpoints, and ultimately drawing users deeper into a narrow content spectrum in the YouTube video recommendation network. Understanding these dynamics is crucial for assessing the societal impacts of content dissemination in digital environments.

In our study, we identified focal structures in YouTube video recommendation networks, represented as a social network $G = (V, E)$, where V is the set of vertices (videos) and E is the set of edges (recommendation links between videos). Focal structures are defined as subgraphs within this network. Formally, focal structures are defined by $F = \{G'\}$, where $G' = (V', E')$, such that $V' \subseteq V$ and $E' \subseteq E$. For all i and j, $i \neq j$, $G_i \in F$ and $G_j \in F$, such that no two focal structures can subsume each other, or $G_i \nsubseteq G_j$ and $G_j \nsubseteq G_i$. [4]. This ensures that the focal structures identified in the network are unique and non-overlapping, each representing a key distinct set of nodes (videos) and edges (recommendations).

D. Network Resiliency Assessment

To identify the key focal structures in our study, we first removed each focal structure from the pruned graph and analyzed the resulting cluster formation to assess network fragmentation. It can be posited that a higher number of clusters indicates greater fragmentation, which suggests that the removed focal structure played a more critical role in maintaining the overall cohesion of the network. In other words, these focal structures act as critical connectors within the overall network. When these key structures are removed, the network's cohesiveness decreases, demonstrating their influence to the overall network [30]. This method emphasizes the structural significance of focal structures and provides insight into how they support content flow throughout the recommendation system and preserve network integrity.

E. Topic Modeling with BERTopic

In order to analyze the content and thematic focus within the focal structures, we applied the BERTopic [31] model, which is a powerful tool for topic modeling that can generate interpretable topics from large text datasets. In our analysis, we opted for BERTopic over conventional models such as Latent Dirichlet Allocation (LDA) [32] or Non-negative Matrix Factorization (NMF) due to its enhanced capability to capture semantic nuances and contextual information. This superiority enables BERTopic to generate more refined and granulated topics, thereby providing deeper insights into the underlying data structure. The primary goal of our study was to determine whether high topic uniformity within a focal structure, where one topic dominates beyond a specified threshold, indicated the presence of a content trap or attractor content, leading to more homogeneous content consumption.

Since YouTube video transcripts are often lengthy, we encountered a limitation due to BERTopic's processing constraint. The maximum length limit of tokens is 512 for the BERTopic model. Hence, we split the video transcripts into multiple chunks, each containing fewer than 512 words, ensuring that each chunk remained coherent by splitting at sentence boundaries. As a result, this method helped maintain the consistency and integrity of the information while ensuring that the chunks adhered to the token limit set by BERTopic. Furthermore, it helped reduce noise, leading to more accurate topic representations. Next, we mapped the identified topics to video IDs to analyze the distribution of videos by topic. This approach enabled a comprehensive understanding of thematic coverage within each focal structure, capturing the diversity or concentration of topics.

After that, we defined a threshold to detect content traps. Specifically, a content trap was identified when a particular topic appeared in more than 50% of the videos within a focal structure. This threshold was chosen based on the idea that a concentration of similar topics across a significant portion of the videos within a focal structure might limit the diversity of content, thereby trapping users in a narrow content trap. In other words, by limiting exposure to varied content, these tightly connected groups of videos increase the likelihood of users being repeatedly nudged toward the same themes, thereby acting as attractor content. The threshold can be formally represented as

$$T = \frac{n_{\text{topic}}}{n_{\text{total}}} > 0.5 \tag{1}$$

where

- T is the threshold for identifying a content trap within a focal structure,
- n_{topic} is the number of videos in the focal structure that share a specific topic, and
- n_{total} is the total number of videos within the focal structure.

We classified the focal structure with a content trap if the proportion T exceeded 0.5. This threshold helped identify clusters where the recommendation algorithm disproportionately favored a single topic that could result in a lack of content diversity. Consequently, this could lead to users continuously engaging with similar or attractor videos over extended periods. Identifying such focal structures can highlight specific areas within YouTube's recommendation system that may contribute to content traps, allowing us to further investigate the implications for user experience and content engagement patterns. Therefore, this methodology provided a systematic way to analyze the topic uniformity in a focal structure within the recommendation system, offering insights into how the role of YouTube's recommendation algorithm may influence content diversity and contribute to content traps.

F. Divergence Metrics

We utilized two key divergence metrics to quantify the similarity between the distributions of topics within focal structures, namely, Kullback-Leibler (KL) Divergence [33] and Jensen-Shannon (JS) Divergence [34]. These metrics provide insights into the diversity and uniformity of topics, which are critical in identifying content traps.

1) Kullback-Leibler (KL) Divergence: The Kullback-Leibler (KL) Divergence measures the difference between two probability distributions P and Q. It is defined as

$$D_{KL}(P||Q) = \sum_i P(i) \log \frac{P(i)}{Q(i)} \tag{2}$$

where

- $P(i)$ represents the true topic distribution and
- $Q(i)$ represents the approximated topic distribution.

A low KL Divergence value indicates that the two distributions are similar, suggesting a limited diversity in the topics present within the focal structure that may signal a content trap.

2) Jensen-Shannon (JS) Divergence: The Jensen-Shannon (JS) Divergence is a symmetrized and smoothed version of KL Divergence. It is computed as

$$D_{JS}(P||Q) = \frac{1}{2}\left(D_{KL}(P||M) + D_{KL}(Q||M)\right) \tag{3}$$

where M is the average distribution

$$M = \frac{1}{2}(P + Q) \tag{4}$$

The JS Divergence is bounded between 0 and 1, making it more stable for comparing distributions. A lower JS Divergence value indicates a lack of topic diversity and the existence of content traps.

In our study, KL Divergence and JS Divergence were employed in order to assess the extent of topic uniformity within focal structures. High values of KL Divergence or JS Divergence suggested a diverse range of topics, indicating the absence of a content trap. In contrast, low KL and JS Divergence values pointed to significant topic uniformity, signaling a potential content trap and providing a quantitative approach to evaluating content traps. Additionally, these approaches offered valuable insights into YouTube's recommendation algorithm dynamics and its role in shaping content diversity.

IV. Experimental Results and Analysis

This section focuses on key factors influencing content traps within YouTube's recommendation network. We start with Focal Structure Analysis and network fragmentation, which examine focal structures and their effects on network resiliency. Next, we explore topical uniformness and content trap identification, focusing on how topic dominance can lead to content homogeneity. The role of divergence metrics is then assessed in identifying content traps through KL and JS Divergence. Lastly, we investigate user engagement's role in reinforcing content homogeneity, highlighting how interactions influence content consumption patterns.

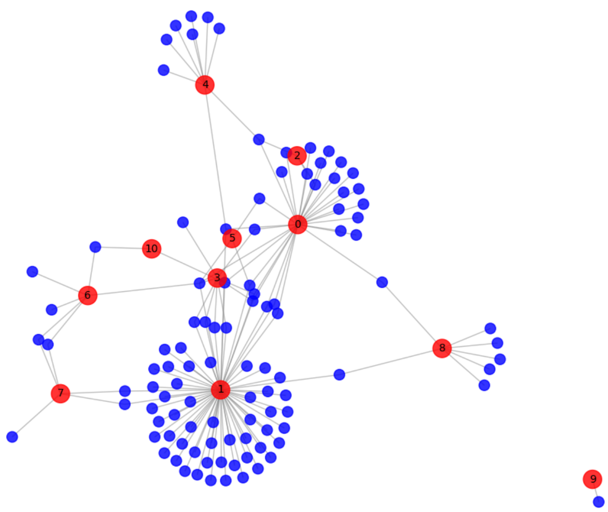

Figure 3: Network visualization of topics and associated video IDs, where red nodes represent topics, blue nodes represent video IDs, and edges denote the association between a video and its topic within the focal structure number 3 of China-Uyghur conflict network data.

A. Focal Structure Analysis and Network Fragmentation

Focal Structure Analysis (FSA) is a social network analysis method that aims to find key sets of individuals rather than a set of key individuals within a social network. The goal of this analysis is to identify structures that can exert influence over the maximum number of entities across various parts and communities within social networks. As part of our research, we first applied the focal structure analysis method to our overall pruned recommendation networks. This analysis identified 105 focal structures within the China–Uyghur conflict dataset and 141 within the Cheng Ho propaganda dataset. After that, we analyzed the network fragmentation resulting from removing each focal structure to identify the key groups in the social network. By assessing the number of clusters formed after the removal, we extrapolated the role of each focal structure in maintaining the network's cohesion. When a focal structure is removed, a greater number of resulting clusters indicates a more significant disruption to the network. This suggests that the removed structure played a crucial role in maintaining connectivity and content flow. Therefore, we were able to identify the most important focal structures for our study. The top 5 of such focal structures for each dataset are listed in Table III.

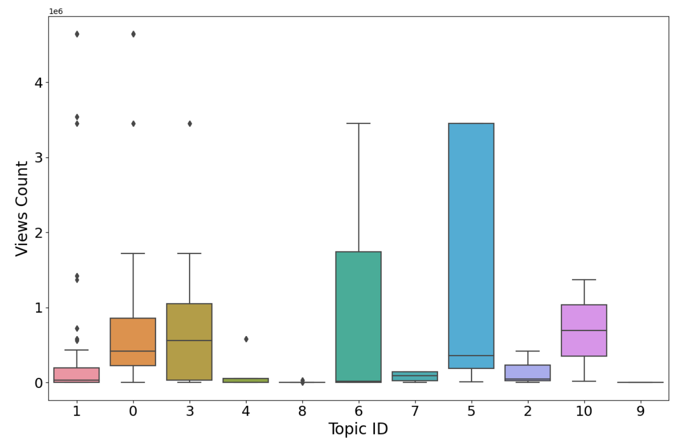

Figure 4: View count distribution across sorted topics based on topic uniformity for China-Uyghur data for focal structure 3. Topic 1 has the highest uniformity but low median views. Significant outliers indicate certain videos received disproportionately high attention, suggesting algorithmic bias and reinforcing the content trap.

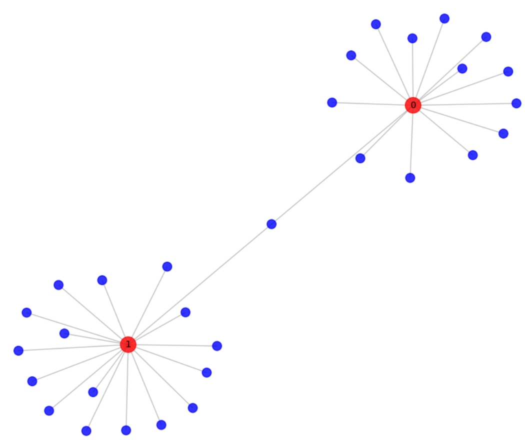

Figure 5: Network visualization of topics and associated video IDs, where red nodes represent topics, blue nodes represent video IDs, and edges denote the association between a video and its topic within the focal structure number 9 of China-Uyghur conflict network data.

Table III: Key metrics for focal structures in the recommendation networks, including size, dominant topic, topic uniformity, and divergence scores. Structures with uniformity above 50% are flagged as potential content traps.

Datasets	Focal Structure (FS)	No. Videos in FS	No. of Videos in Dominant Topic	No. of Clusters	Topic Uniformity	Content Trap	KL Divergence	JS Divergence
China-Uyghur Dataset	3	105	64	185	61%	YES	0.680	0.158
	9	30	17	44	57%	YES	0.004	0.001
	1	25	15	41	60%	YES	0.012	0.003
	102	13	7	31	54%	YES	0.067	0.234
	101	13	5	28	38%	NO	0.154	0.043
Cheng Ho Propaganda Dataset	13	68	45	134	66%	YES	0.700	0.166
	1	52	21	96	40%	NO	0.770	0.174
	6	28	22	81	79%	YES	0.446	0.120
	22	29	16	74	55%	YES	0.094	0.025
	3	28	18	74	64%	YES	0.277	0.069

B. Topic Uniformity and Content Trap Identification in Focal Structures

In this section, we delve deeper into analyzing the thematic focus within the identified key focal structures by measuring the topic uniformity. After identifying the top key focal structures from our networks, we applied the BERTopic model to extract interpretable topics from the video transcripts. Next, for each focal structure, we extracted the number of videos associated with each topic. Topic uniformity was then quantified by calculating the dominance of a particular topic within a set of videos, such as those within a focal structure. This measure reflects the proportion of content assigned to the dominant topic, which indicates the extent to which a focal structure is centered around a single or attractor topic. By setting up a threshold for topic uniformity, we identified content traps. A focal structure was classified as having a content trap if a single or attractor topic appeared in more than 50% of the videos within that structure. Table III presents the total number of videos in each focal structure, the number of videos in the dominant topic associated with each focal structure, topic uniformity, and whether the structure was identified as a content trap. This methodology provides valuable insight into how homogeneous topic concentration within focal structures can lead to a lack of content diversity, potentially trapping users in a narrow content trap and addressing our research question **RQ1**. In addition, to support these findings, Figures 3 and 5 illustrate the visualization of topics alongside their corresponding video IDs for two such focal groups within the China Uyghur conflict dataset which are also the two topmost focal structures that identified as content trap. Similarly, the top two focal structures that are identified as content traps from Table III are shown in Figures 7 and 9, which present comparable visualizations for the Cheng Ho propaganda network, where all the figures highlight a clear pattern of topics clustering together, suggesting the presence of attractor content. As a result, this phenomenon highlights a significant reduction in thematic diversity, where discourse becomes increasingly concentrated around a limited set of ideas or contents.

C. Divergence Metrics and Their Role in Identifying Content Traps

When evaluating the topical consistency within focal structures for our study, the Jensen-Shannon (JS) and Kullback-Leibler (KL) divergence metrics played a significant role.

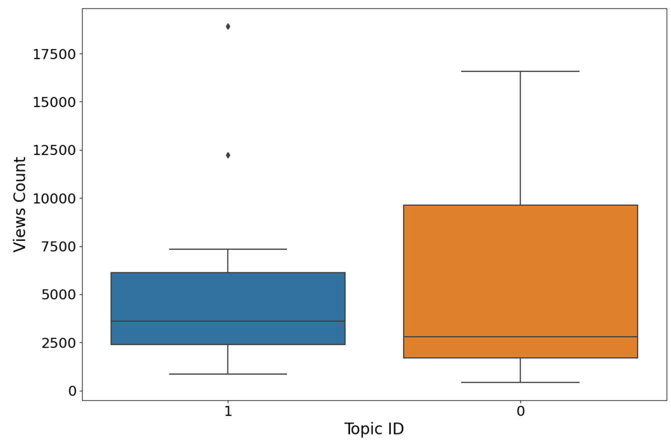

Figure 6: View count distribution across sorted topics based on topic uniformity for China-Uyghur data for focal structure 9. Topic 1 has the highest uniformity and higher median view count. Significant outliers indicate certain videos received disproportionately high attention, suggesting algorithmic bias and reinforcing the content trap.

These metrics helped us identify potential content traps within YouTube's recommendation system by measuring the distribution of topics within a focal structure. When a focal structure exhibits low divergence values, it suggests that the recommendation algorithm primarily favors a limited set of topical themes, reinforcing homogeneous content consumption. In contrast, higher divergence values indicate a broader distribution of topics with less topical uniformity, suggesting a more diverse range of content and thereby reducing the likelihood of a content trap. By analyzing these divergence metrics as shown in Table III, we identified content traps within focal structures. For instance, Focal Structure 3 for the China–Uyghur conflict dataset stands out with a relatively higher KL Divergence value compared to other focal structures, but it also exhibits a much lower JS Divergence. This combination suggests that although there may be some degree of variability in topic distribution as indicated by the KL Divergence, the overall diversity of topics within this structure remained limited, as evidenced by the low JS Divergence value. This leads to a content trap scenario, where a single or attractor topic dominates the recommendations, restricting the variety of content available to users.

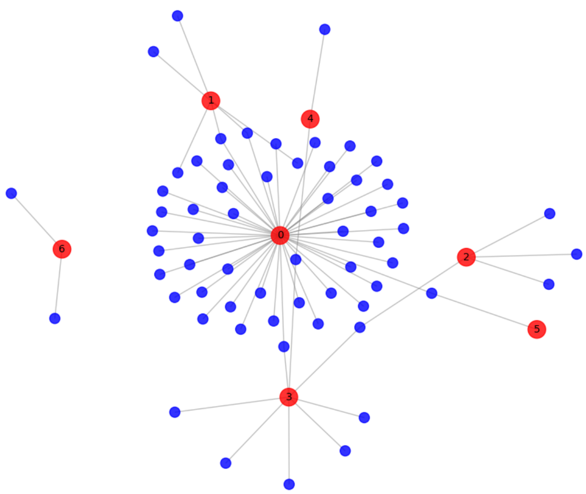

Figure 7: Network visualization of topics and associated video IDs, where red nodes represent topics, blue nodes represent video IDs, and edges denote the association between a video and its topic within the focal structure number 13 of Chengo Ho propaganda.

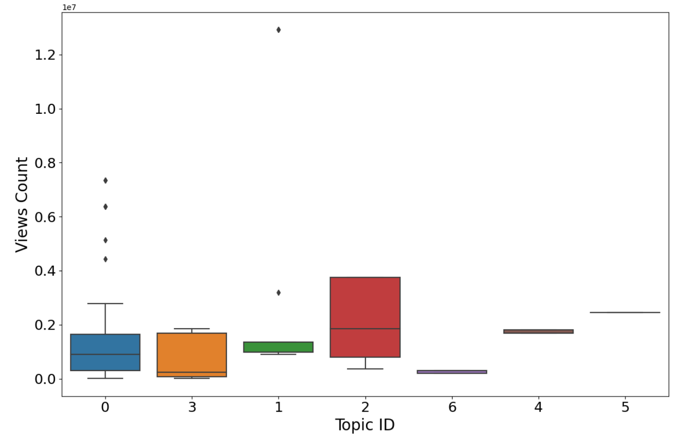

Figure 8: View count distribution across sorted topics based on topic uniformity for Cheng Ho data for focal structure 13. Topic 0 has the highest uniformity but low median views. Significant outliers indicate certain videos received disproportionately high attention, suggesting algorithmic bias and reinforcing the content trap.

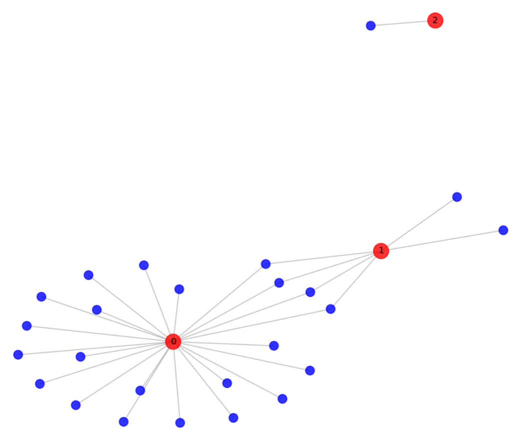

Figure 9: Network visualization of topics and associated video IDs, where red nodes represent topics, blue nodes represent video IDs, and edges denote the association between a video and its topic within the focal structure number 6 of Chengo Ho propaganda.

However, not all focal structures with high divergence scores resulted in content traps. For example, Focal Structure 1 in the Chengo Ho propaganda dataset exhibited both high KL and JS Divergence values, indicating that it still maintained internal topical diversity. This suggests that users exposed to this focal structure encountered a broader range of content rather than being constrained to a single dominant theme. As a result, this type of structure does not reinforce homogeneous content consumption and thus reduces the likelihood of content entrapment. By incorporating both cases, our analysis provides a nuanced understanding of how topical consistency influences content diversity within focal structures, thereby addressing our research question **RQ2**.

D. Role of User Engagement in Reinforcing Homogeneous Content

In our analysis, topics with higher uniformity exhibited notable outliers in user engagement metrics such as view count. In other words, while topics with high topic uniformity reflected a concentration around a single theme, the engagement distribution unveiled a more complex pattern, with some videos receiving disproportionately high attention. User engagement from the highly uniform topics showed disproportionately high engagement as well as surpassing the median engagement level in other topics for certain focal structures. However, the disproportionate high engagement persisted in most of the focal structures' high uniformity topics. This highlights how high-engagement videos reinforce homogeneous content consumption by amplifying specific content through the recommendation algorithm, further promoting uniformity within topics. Figures 4 and 6 illustrate that topics exhibiting the highest levels of topic uniformity also showed a greater number of outliers compared to other topics within the focal structure for the China–Uyghur conflict network, and Figures 8 and 10 portray a similar kind of pattern in the Cheng Ho propaganda network as well. This dynamic underscores how engagement metrics and topic uniformity can significantly boost content traps, where diversity is limited and users' exposure to new or varied content is constrained and thus addressing our research question **RQ3**.

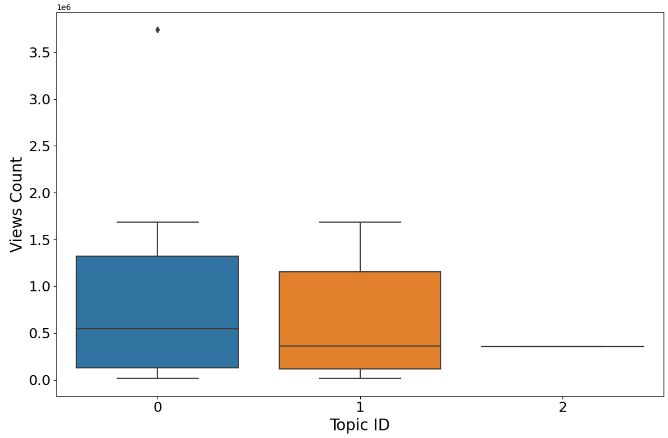

Figure 10: View count distribution across sorted topics based on topic uniformity for Cheng Ho data for focal structure 6. Topic 0 has the highest uniformity and higher median view count. Significant outliers indicate certain videos received disproportionately high attention, suggesting algorithmic bias and reinforcing the content trap.

V. Conclusions and Future Work

This study investigated how YouTube's recommendation network can give rise to content traps by jointly examining their structural, semantic, and behavioral properties. Our analysis shows that focal structures represent highly cohesive sub-networks that play a central role in maintaining the overall connectivity of the recommendation graph. Answering **RQ1**, we found that the removal of these structures causes substantial network fragmentation, indicating their importance as structural attractor groups. Addressing **RQ2**, topic modeling and divergence-based analyses demonstrated that several of these attractors exhibit pronounced topical uniformity, reflecting limited thematic diversity within the content they promote. Finally, in response to **RQ3**, engagement patterns revealed that videos within the most uniform topics receive disproportionately high user attention, suggesting that user behavior and algorithmic amplification jointly reinforce these content traps.

The lack of direct comparison with state-of-the-art community detection and influential node detection methods may seem to be a limitation of this study. However, prior research has shown that focal structure analysis outperforms traditional community and influential node discovery methods [4]. Given that, it can be posited that focal structure analysis offers better insights into content traps than traditional methods. Future work could expand this analysis across other online platforms which could further provide a deeper understanding of the dynamics shaping content flow and user engagement.

Acknowledgements

This research is funded in part by the U.S. National Science Foundation (OIA-1946391, OIA-1920920), U.S. Office of the Under Secretary of Defense for Research and Engineering (FA9550-22-1-0332), U.S. Army Research Office (W911NF-23-1-0011, W911NF-24-1-0078, W911NF-25-1-0147), U.S. Office of Naval Research (N00014-21-1-2121, N00014-21-1-2765, N00014-22-1-2318), U.S. Air Force Research Laboratory, U.S. Defense Advanced Research Projects Agency, the Australian Department of Defense Strategic Policy Grants Program, Arkansas Research Alliance, the Jerry L. Maulden/Entergy Endowment, and the Donaghey Foundation at the University of Arkansas at Little Rock. 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] M. Bhuiyan and N. Agarwal, "Identification and characterization of content traps in youtube recommendation network", in *eKNOW 2025, The Seventeenth International Conference on Information, Process, and Knowledge Management*, IARIA, 2025, pp. 59–64.

[2] C. Agency, *Youtube statistics, https://www.charleagency.com/articles/youtube-statistics/*, Accessed: April 10, 2025, 2024.

[3] A. Gallagher, L. Cooper, R. Bhatnagar, and C. Gatewood, *Pulling back the curtain: An exploration of youtube's recommendation algorithm*, 2024.

[4] M. Alassad, M. N. Hussain, and N. Agarwal, "Comprehensive decomposition optimization method for locating key sets of commenters spreading conspiracy theory in complex social networks", *Central European Journal of Operations Research*, vol. 30, no. 1, pp. 367–394, 2022.

[5] M. A. Brown *et al.*, "Echo chambers, rabbit holes, and algorithmic bias: How youtube recommends content to real users", *Available at SSRN 4114905*, 2022.

[6] M. M. I. Bhuiyan and N. Agarwal, "Detecting algorithmic homophily in recommendation graphs via weighted topic distribution", *2025 IEEE 37th International Conference on Tools with Artificial Intelligence (ICTAI)*, 2025.

[7] J. M. Kleinberg, "Authoritative sources in a hyperlinked environment", *Journal of the ACM (JACM)*, vol. 46, no. 5, pp. 604–632, 1999.

[8] S. Brin and L. Page, "The anatomy of a large-scale hypertextual web search engine", *Computer Networks and ISDN Systems*, vol. 30, no. 1-7, pp. 107–117, 1998.

[9] P. Bedi and C. Sharma, "Community detection in social networks", *Wiley interdisciplinary reviews: Data Mining and Knowledge Discovery*, vol. 6, no. 3, pp. 115–135, 2016.

[10] F. Sen, R. T. Wigand, N. Agarwal, D. Mahata, and H. Bisgin, "Identifying focal patterns in social networks", in *2012 Fourth International Conference on Computational Aspects of Social Networks (CASoN)*, IEEE, 2012, pp. 105–108.

[11] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, "Fast unfolding of communities in large networks", *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2008, no. 10, P10008, 2008.

[12] F. Şen, R. T. Wigand, N. Agarwal, M. Mete, and R. Kasprzyk, "Focal structure analysis in large biological networks", in *3rd International Conference on Environment, Energy and Biotechnology (ICEEB 2014)*, 2014.

[13] F. Şen, R. Wigand, N. Agarwal, S. Tokdemir, and R. Kasprzyk, "Focal structures analysis: Identifying influential sets of individuals in a social network", *Social Network Analysis and Mining*, vol. 6, pp. 1–22, 2016.

[14] M. J. Alenazi and J. P. Sterbenz, "Comprehensive comparison and accuracy of graph metrics in predicting network resilience", in *2015 11th International Conference on the Design of Reliable Communication Networks (DRCN)*, IEEE, 2015, pp. 157–164.

[15] E. Pariser, *The filter bubble: How the new personalized web is changing what we read and how we think*. Penguin, 2011.

[16] A. Bechmann and K. L. Nielbo, "Are we exposed to the same "news" in the news feed? an empirical analysis of filter bubbles as information similarity for danish facebook users", *Digital Journalism*, vol. 6, no. 8, pp. 990–1002, 2018.

[17] G. M. Lunardi, G. M. Machado, V. Maran, and J. P. M. de Oliveira, "A metric for filter bubble measurement in recommender algorithms considering the news domain", *Applied Soft Computing*, vol. 97, p. 106 771, 2020.

[18] F. Masrour, T. Wilson, H. Yan, P.-N. Tan, and A. Esfahanian, "Bursting the filter bubble: Fairness-aware network link prediction", in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, 2020, pp. 841–848.

[19] L. Burbach, P. Halbach, M. Ziefle, and A. Calero Valdez, "Bubble trouble: Strategies against filter bubbles in online social networks", in *Digital Human Modeling and Applications in Health, Safety, Ergonomics and Risk Management. Healthcare Applications: 10th International Conference, DHM 2019, Held as Part of the 21st HCI International Conference, HCII 2019, Orlando, FL, USA, July 26–31, 2019, Proceedings, Part II 21*, Springer, 2019, pp. 441–456.

[20] A. Amrollahi, "A conceptual tool to eliminate filter bubbles in social networks", *Australasian Journal of Information Systems*, vol. 25, pp. 1–16, 2021.

[21] R. Finlay, "The voyages of zheng he: Ideology, state power, and maritime trade in ming china", *Journal of the Historical Society*, vol. 8, no. 3, pp. 327–347, 2008.

[22] A. M. Dwyer, "The xinjiang conflict: Uyghur identity, language policy, and political discourse", *Policy Studies*, vol. 15, 2005.

[23] R. Hasmath, "What explains the rise of majority–minority tensions and conflict in xinjiang?", *Central Asian Survey*, vol. 38, no. 1, pp. 46–60, 2019.

[24] M. C. Cakmak, O. Okeke, U. Onyepunuka, B. Spann, and N. Agarwal, "Investigating bias in youtube recommendations: Emotion, morality, and network dynamics in china-uyghur content", in *International Conference on Complex Networks and Their Applications*, Springer, 2023, pp. 351–362.

[25] M. C. Cakmak, N. Agarwal, and R. Oni, "The bias beneath: Analyzing drift in youtube's algorithmic recommendations", *Social Network Analysis and Mining*, vol. 14, no. 1, p. 171, 2024.

[26] Google Developers, *YouTube Data API*, https://developers.google.com/youtube/v3, Accessed: April. 10, 2025, 2024.

[27] M. C. Cakmak and N. Agarwal, "High-speed transcript collection on multimedia platforms: Advancing social media research through parallel processing", in *2024 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, 2024, pp. 857–860. DOI: 10.1109/IPDPSW63119.2024.00153.

[28] G. Zhang *et al.*, "Gder: Safeguarding efficiency, balancing, and robustness via prototypical graph pruning", *Advances in Neural Information Processing Systems*, vol. 37, pp. 50 285–50 312, 2024.

[29] J. Li *et al.*, "Less can be more: Unsupervised graph pruning for large-scale dynamic graphs", *arXiv preprint arXiv:2305.10673*, 2023.

[30] M. Bhuiyan, S. Shajari, and N. Agarwal, "Resilience and node impact assessment in youtube commenter networks leveraging focal structure analysis", *The Eleventh International Conference on Human and Social Analytics (HUSO 2025)*, 2025.

[31] M. Grootendorst, "Bertopic: Neural topic modeling with a class-based tf-idf procedure", *arXiv preprint arXiv:2203.05794*, 2022.

[32] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation", *Journal of Machine Learning Research*, vol. 3, no. Jan, pp. 993–1022, 2003.

[33] S. Kullback and R. A. Leibler, "On information and sufficiency", *The Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79–86, 1951.

[34] J. Lin, "Divergence measures based on the shannon entropy", *IEEE Transactions on Information Theory*, vol. 37, no. 1, pp. 145–151, 1991.