Revealing Structural Influence in AI-Driven Recommendation Systems: A Weighted Focal Structure Analysis of YouTube Networks

Nitin Agarwal[1,2], **Bishwa Prakash Subedi**[1]
[1]COSMOS Research Center, University of Arkansas-Little Rock, USA
[2]International Computer Science Institute, University of California, Berkeley, USA
nxagarwal@ualr.edu, bsubedi@ualr.edu

Abstract

AI-driven recommendation systems increasingly influence how people discover and consume online content. Understanding how these systems shape user navigation has therefore become an important research problem. YouTube's recommendation algorithm, which drives nearly 70% of watch time on the platform, acts as a continuous intermediary between users and the content they encounter. In this study, we apply Weighted Focal Structure Analysis (WFSA) to YouTube recommendation networks, introducing edge weights based on recommendation rank and depth to reflect how recommendation hierarchy influences network structure. Using datasets drawn from the 2024 Taiwan presidential election and the 2025 U.S. tariff expansion, we construct weighted recommendation graphs and extract focal structures using both FSA and five WFSA variants. Our results show that WFSA consistently identifies smaller, more tightly connected, and structurally influential groups than FSA does. Network resiliency experiments show that removing WFSA focal structures causes greater network fragmentation and connectivity loss. Among the evaluated weighting schemes, WFSA4 (linear rank decay, inverse depth decay with count-weighted aggregation) most effectively captures structurally critical groups. These findings suggest that recommendation hierarchy plays an important role in shaping which groups of content become influential within recommendation networks.

Introduction

Across the modern web, AI-driven recommendation systems have become primary intermediaries between people and the information they encounter. Rather than users actively seeking out content, algorithmically curated platforms increasingly decide what appears next, ultimately shaping attention, directing navigation, and playing a significant role in how people form views and consume news. YouTube, which operates in more than 100 countries and supports over 80 languages, is one of the most widely used platforms of this kind and the world's second-most-visited social media site (WordStream 2026). Studies estimate that nearly 70% of YouTube watch time is driven by automated recommendations (Institute for Strategic Dialogue 2024). This means the recommendation algorithm is not a neutral suggestion mechanism; it is an active participant in the human experience of the platform, continuously making decisions about which content pathways users are likely to follow.

Since recommendations connect videos through successive viewing paths, the system can be modeled as a recommendation graph capturing content linkage and information flow. These algorithms may also produce structural patterns that repeatedly guide users toward similar content, a phenomenon known as content traps (Bhuiyan and Agarwal 2025), making it important to study how such structures shape AI-human interaction.

Prior work (Bhuiyan and Agarwal 2025) addressed this by applying Focal Structure Analysis (FSA) to extract influential groups from the recommendation networks and examine their thematic consistency. However, the traditional FSA framework assumes an unweighted network and treats all the recommendation links equally. In YouTube's recommendation hierarchy, this assumption does not account for behavioral signals. At the same recommendation depth, videos near the top of the list attract more user attention and are more likely to be clicked, while those lower in the list are less likely to be selected. In addition, videos closer to the seed video in the recommendation chain are more likely to be explored by users, whereas those farther away tend to receive less attention. Ignoring these differences may lead to an incomplete understanding of how algorithmic structure shapes human navigation behavior.

To address this limitation, we extend the application of Weighted Focal Structure Analysis (WFSA) to YouTube recommendation networks by introducing edge weights derived from recommendation rank and depth, signals specific to the hierarchical logic of AI-driven recommendation systems that have not been explored in prior WFSA work. Unlike traditional approaches that treat all connections equally, WFSA explicitly embeds the hierarchical logic of AI-driven recommendation systems into network analysis. By doing so, it enables the identification of influential groups shaped not only by structural connectivity but also by algorithmic content prioritization. This work bridges the gap between network-based influence analysis and the hierarchical structure of AI recommendation systems, providing a more realistic representation of how influence emerges in algorithmically curated environments.

Using this framework, we systematically compare focal structures extracted by FSA and multiple WFSA variants to

---

understand how weighting affects the identification of influential groups and the structural resilience of the recommendation networks. This study investigates the following research questions:

- RQ1: Which focal structures (or focal sets), identified by FSA or WFSA, are more structurally important in the network?
- RQ2: Which approach, FSA or WFSA, leads to greater network disruption when its focal structures are removed?
- RQ3: Among multiple WFSA schemes, which weighting scheme performs the best?

Understanding these structural patterns is important for studying how recommendation systems guide user navigation and shape information pathways in algorithmically curated platforms. Through the proposed WFSA framework and comparative analysis, this study provides insights into how recommendation hierarchy influences the identification of structurally influential groups and how these groups affect the connectivity of recommendation networks. The remainder of this paper is organized as follows. The literature review summarizes relevant prior research, followed by the background and data collection sections, which describe the discourses represented in the datasets and outline the data collection process. The methodology and results sections then present the research approach and key findings. Finally, the paper discusses the implications of these findings and concludes with limitations and directions for future work.

Literature Review

Early research on influence in networks primarily focused on identifying important individual nodes. Classical centrality measures such as degree, betweenness, and closeness centrality have been widely used to quantify the structural importance of nodes within a network (Freeman 1978). These metrics evaluate how well a node is connected, how frequently it lies on shortest paths between other nodes, and how close it is to the rest of the network. With the growth of large-scale online networks, recursive ranking algorithms such as PageRank (Brin and Page 1998) and HITS (Kleinberg 1999) became widely adopted to capture influence based on link structures, modeling importance as a recursive process in which a node's significance depends on the significance of nodes linking to it.

Although node-level measures are useful for identifying prominent actors, influence in many social systems often emerges through groups of coordinated nodes rather than isolated individuals. To capture this group-level structure, researchers introduced community detection methods that aim to identify clusters of nodes with dense internal connections and relatively sparse external links (Bedi and Sharma 2016). Methods such as modularity optimization and the Louvain algorithm have been widely used to reveal community structure in large networks. However, communities are often large and heterogeneous, and therefore may not always represent the smaller sets of actors that drive influence or coordination within a network.

To address this limitation, Focal Structure Analysis (FSA) was proposed to identify small, structurally cohesive groups of nodes that exert collective influence within a network (Sen et al. 2012; Şen et al. 2016). Unlike community detection approaches that identify broad clusters, FSA focuses on compact groups that are both internally cohesive and strategically positioned within the network. These focal structures capture coordinated sets of nodes that can significantly influence the network's connectivity and information flow. Prior research has applied FSA to various types of social networks to identify influential subgroups and examine how their removal affects network stability.

More recently, FSA has been applied to recommendation networks to study content traps in algorithmically curated platforms. Bhuiyan and Agarwal (Bhuiyan and Agarwal 2025) demonstrated how recommendation pathways can guide users toward groups of closely related videos that repeatedly reinforce thematically similar content, effectively narrowing the diversity of information that users encounter. These findings highlight how recommendation network structures can influence the pathways through which users encounter content and may contribute to the formation of content traps. Researchers studying YouTube's recommendation system have observed similar problems with how recommendations are structured (Ribeiro et al. 2020).

Researchers have also examined broader dynamics of algorithmic recommendations. Modern systems rely on complex ranking architectures guiding users through sequential recommendations (Covington, Adams, and Sargin 2016), and YouTube's system can gradually push users toward specific content types, narrowing information diversity (Cakmak, Agarwal, and Oni 2024). Studies show that ranked presentation disproportionately affects user attention and engagement (Joachims et al. 2007; Craswell et al. 2008), and algorithmic curation has been linked to filter bubbles and narrower information exposure (Pariser 2011; Bakshy, Messing, and Adamic 2015), reinforcing the need to account for hierarchical signals in recommendation network analysis.

Despite these advances, most existing studies apply focal structure analysis using unweighted network representations. In recommendation systems, however, connections between nodes are not uniform. Recommendation links vary in importance based on rank and navigation depth, both of which directly affect the likelihood that a human user will follow a given path. Ignoring these signals may cause analysis to miss the patterns through which AI systems concentrate structural influence.

Recent work has extended FSA to weighted networks under the Weighted Focal Structure Analysis (WFSA) framework. Falade and Agarwal (Falade and Agarwal 2026) introduced WFSA to detect toxic focal structures in social networks by using toxicity-weighted edges, thereby improving the detection of influential groups. Subsequent work further applied WFSA to toxicity propagation modeling (Falade and Agarwal 2025b) and optimized the selection of focal toxic structures using integer programming (Falade and Agarwal 2025a). However, none of these works have applied WFSA to recommendation networks, in which edge weights are derived from hierarchical signals embedded in the recommendation system itself, namely, rank position and navigation depth.

To address this gap, we extend the traditional Focal Structure Analysis (FSA) framework by introducing a Weighted Focal Structure Analysis (WFSA) approach that incorporates edge weights derived from recommendation rank and depth. By systematically comparing focal structures identified using FSA and multiple WFSA variants, this study examines how recommendation hierarchy influences the identification of influential groups within recommendation networks. Therefore, the work provides insight into how hierarchical signals embedded in AI recommendation systems shape the structural properties of recommendation pathways and which content groups become structurally influential.

Background and Data Collection

We provide essential context for the discourses from which the data are collected. The datasets used in this study are drawn from two major political and economic topics that generated substantial online discussion: the 2024 Taiwan presidential election and the expansion of U.S. tariff policies in early 2025. Both events received significant global media coverage and sparked extensive conversations on social media platforms, making them suitable contexts for analyzing recommendation network structures in which algorithmic systems influence how users navigate and encounter information.

2024 Taiwan Presidential Election: The Taiwanese presidential election held on January 13, 2024, was a significant political event that attracted widespread domestic and international attention. The election featured three main candidates representing different political positions regarding Taiwan's governance and its relationship with mainland China. Lai Ching-te of the Democratic Progressive Party (DPP) supported maintaining Taiwan's existing autonomous status, while Hou Yu-ih of the Kuomintang (KMT) advocated closer cross-strait engagement. The election period generated extensive public discussion across online platforms, where political commentary, campaign messaging, and election-related information circulated widely (Gan 2024).

Tariff Expansion 2025: In early 2025, the United States introduced a series of tariff increases on imported goods as part of a broader shift in trade policy. Tariff rates on several categories of products were raised, in some cases reaching approximately 10%–25% (Lawder and Shalal 2025). These policy changes quickly became a subject of public and economic debate, particularly regarding their potential effects on consumer prices, supply chains, and international trade relations. As a result, the topic generated substantial media coverage and online discussions, making it a relevant context for examining recommendation network dynamics.

Data Collection

The data for this study were collected from YouTube across three discourse domains: Tariff (US), Tariff (China), and the

Table 1: Keywords used to identify seed videos

Discourse	Keywords
Tariff China	川普, 關稅, chinatariff, 經濟, 國際稅收, 美國稅收制度, 川普政府, 關稅政策, 相互關稅, 자동차관세, 트럼프관세, 관세영향, 川普總統, 川普的四年是一場惡夢, 川普關稅, 美元, 관세, 트럼프, 王牌, ChinaSanctions, ChinaTariff, ChinaTariffHike, ChinaTariffs, ChinaTradeWar, ChinaUSAtradeWar, ChinaUSTradeWar, ChinaUSTradesWar, ChinaWarning, ChinaWarns, Chinatariffs, ChineseTariffs, chinatariffnews, chinatariffs, chinatariffsonus, chinesetariffs
Tariff US	TradeWar, TrumpTariffs, TariffImpact, TradeWars, TradeNegotiations, USChinaTrade, PriceHikes, TariffWar, TariffThreat, EconomicCrisis, EconomicImpact, LiberationDay, AmericaFirst, Tariff surcharge, Sticker shock, Price creep, Re-ticketing, Fashion tariff, Electronics cost, Food prices, Shipping fees, Supply chain whipsawing, Margin compression, Landed costs, Empty shelves, Operational uncertainty, Hidden tax, Economic coercion, Trade war chaos, Regressive tax effects, Constitutional overreach, and Economic weapons
Taiwan	DemocraticProgressiveParty, DPP, KoWenje, KuomintangParty, kuomingtang, Kuomintang, LaiChingte, LegislativeYuan, Mygopen, NewPowerParty, prayfortaiwan, SpeakOutDontFight, TaiwanElection, taiwanelection2024, TaiwanElections, taiwanisacountry, taiwanisnotapartofchina, taiwanisnotchina, TaiwanPresidentialElection2024, TsaiIng-wen, votefortaiwan, whiteterror, WilliamLai, YouXikun, 民進, 蔣介石, 中國國民黨, 侯友宜, 台灣立法院選, 台灣總統選, 台灣大選, 台灣民黨, 台灣總統大選2024, 國民黨, 民主進, 白色恐怖, 立委選, 總統大選, 蔡英文, 蕭美琴, 賴德, 選前之夜

Taiwan presidential elections. For each discourse, we constructed a recommendation network starting from a set of seed videos. To ensure broad and representative coverage of each topic, seed videos were selected using three sampling strategies: viral videos, non-viral videos, and three randomly selected sample sets. As a result, five datasets were constructed for each discourse domain: viral, non-viral, sample 1, sample 2, and sample 3.

Seed videos were identified using keyword lists carefully selected by reviewing the newspaper articles covering these discourses (Boak and Sanchez 2025; Lawder and Shalal 2025; Gan 2024). The keywords are presented in Table 1. These keywords were used to retrieve videos that were relevant to each discourse while minimizing unrelated content. For each discourse domain, 50 seed videos were selected. Starting from each seed video, we collected YouTube recommendations using a breadth-first crawling process. The crawler expanded the recommendation network up to 5 levels (hops), retrieving at most 5 recommendations from each video at each expansion level. Limiting the crawl to five hops allowed us to capture the hierarchical structure of the recommendation network while maintaining computational feasibility.

The data collection process yielded 15 recommendation networks across the three discourse domains. Each record in the resulting dataset contains the root video identifier, the recommended video identifier, the depth level of the recommendation within the crawl, and the rank position of the recommended video at that depth. No personally identifiable information (PII) was collected during this process.

Table 2: Node and edge counts of the 15 recommendation graphs constructed across three discourse domains.

S.N	Dataset	Nodes	Edges
1	China non-viral sample	8,198	14,520
2	China viral sample	11,880	19,504
3	China sample 1	6,552	12,659
4	China sample 2	7,692	14,894
5	China sample 3	9,379	17,638
6	Taiwan non-viral sample	7,542	13,433
7	Taiwan viral sample	6,439	13,436
8	Taiwan sample 1	7,850	14,816
9	Taiwan sample 2	9,653	17,983
10	Taiwan sample 3	9,480	17,541
11	US non-viral sample	5,313	9,442
12	US viral sample	10,665	20,513
13	US sample 1	8,644	16,956
14	US sample 2	9,291	19,981
15	US sample 3	9,388	19,695

Methodology

This section describes the methodology followed in this study. First, data were collected from YouTube and used to construct recommendation networks for the selected discourse domains. After building the recommendation graphs, edge-weighting schemes based on recommendation rank and depth were applied to generate weighted networks. Next, focal structures (FS) were extracted from both the unweighted and weighted graphs using FSA and WFSA approaches. The structural properties of the focal structures identified by each method were then analyzed and compared. Finally, a network resiliency assessment (Bhuiyan, Shajari, and Agarwal 2025) was conducted to evaluate the impact of removing these focal structures on the connectivity and stability of the recommendation networks. Figure 1 represents the overall research methodology followed.

Weighted Recommendation Graph Construction

The construction of the unweighted recommendation graph is relatively straightforward: duplicate edges between the same source node (A) and target node (B) are removed to maintain a clean structure. However, treating all recommendation links equally overlooks important signals that shape users' experience of the platform. In practice, recommendation systems present content in a ranked and structured manner, where factors such as position and depth affect user exposure and engagement. Ignoring these factors can lead to an incomplete representation of the recommendation process.

To address this, weighted recommendation graphs are constructed by assigning edge weights based on the mechanisms described in this section. This weighted representation forms the basis for the proposed Weighted Focal Structure Analysis (WFSA), enabling a more realistic modeling of recommendation pathways and their structural influence. Table 2 summarizes the resulting recommendation graphs, including the number of nodes and edges in each network, providing insight into their overall size.

In YouTube's recommendation system, both rank and depth carry meaningful behavioral information. Recommendations appearing at higher ranks (e.g., rank = 1) are typically more visible to users and therefore more likely to be clicked. Recommendations appearing farther from the seed video correspond to longer navigation paths that users are unlikely to follow. Treating all edges equally, therefore, ignores the hierarchical nature of the recommendation process. Because the recommendation system orders videos and structures recommendation chains, videos placed higher in the list and closer to the seed video tend to receive greater user attention, which in turn affects which content becomes structurally influential within the recommendation network.

To incorporate these signals, we extend the recommendation graph by assigning edge weights based on rank and depth. The weighting scheme is designed to approximate realistic user navigation behavior:

- Users are more likely to click recommendations appearing at higher ranks (rank 1, rank 2, rank 3, etc.).
- The visibility of recommended videos decreases as the depth in the recommendation chain increases, since users rarely navigate far along multi-hop recommendation paths.

By incorporating both rank and depth into the weighting scheme, the resulting weighted graph better reflects the

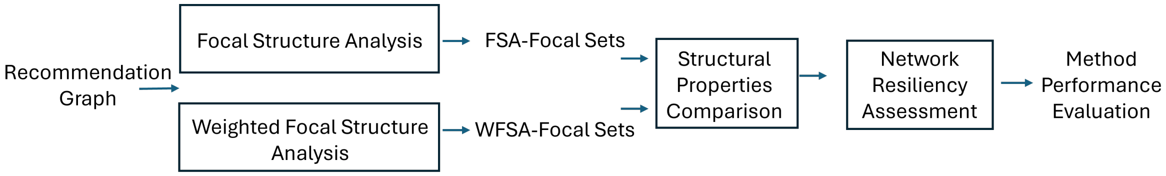

Figure 1: Methodology pipeline for FSA and WFSA-based focal structure extraction and evaluation.

structural importance of recommendation pathways and provides a more realistic representation of how users traverse the AI-driven recommendation system.

Rather than assuming a single correct model of user attention decay, we explore multiple functional forms to examine how different approximations of hierarchical navigation behavior affect the identification of focal structure. This exploration is motivated by position-bias research demonstrating that user attention decreases with rank position (Joachims et al. 2007; Craswell et al. 2008), while the specific functional forms, exponential and linear for rank, exponential and inverse for depth, represent a range of common decay behaviors from aggressive to gradual. The relative performance of these schemes is then evaluated empirically, with the results determining which approximation most effectively captures structurally influential groups. Specifically, we explored several weighting strategies based on different combinations of rank and depth decay functions.

1. **Exponential–Exponential Decay**

   In this approach, both rank and depth importance decrease exponentially.

   $$R_{\text{wt}} = e^{-(\text{rank}-1)} \quad (1) \qquad D_{\text{wt}} = e^{-(\text{depth}-1)} \quad (2)$$

2. **Exponential–Inverse Decay**

   In this approach, rank importance decreases exponentially, while depth importance decreases inversely.

   $$R_{\text{wt}} = e^{-(\text{rank}-1)} \quad (3) \qquad D_{\text{wt}} = \frac{1}{\text{depth}} \quad (4)$$

3. **Linear–Inverse Decay**

   In this approach, rank importance decreases linearly, while depth importance decays inversely.

   $$R_{\text{wt}} = 1 - \frac{\text{rank}-1}{\text{maxrank}} \quad (5) \qquad D_{\text{wt}} = \frac{1}{\text{depth}} \quad (6)$$

After computing the rank and depth weights, the final edge weight is calculated using two alternative aggregation strategies.

$$w = \frac{\sum (R_{\text{wt}} D_{\text{wt}})}{\sum (R_{\text{wt}})} \quad (7) \qquad w = \frac{\sum (R_{\text{wt}} D_{\text{wt}})}{\text{Total occurrence}} \quad (8)$$

These aggregation strategies allow the weighting process to account for both the strength of individual recommendation links and the frequency with which a recommendation pair appears in the dataset.

FSA

Focal Structure Analysis (FSA) is a social network analysis technique that identifies small, influential groups of nodes, known as focal structures, that exert collective impact within a network (Sen et al. 2012; Şen et al. 2016). In contrast to methods centered on individual centrality scores or large community partitions, FSA isolates compact subsets of nodes that are tightly connected and strategically embedded within the network. In this study, we adopt the bi-level optimization-based FSA framework proposed by Alassad et al., (Alassad, Hussain, and Agarwal 2019). By prioritizing both group-level influence and internal cohesiveness, FSA reveals clusters that function as effective units rather than relying on the power of single actors. These focal structures are particularly useful for studying information spread, behavioral dynamics, and structural weaknesses in social and recommendation networks.

WFSA

WFSA is an extended version of FSA designed for weighted recommendation networks. In this setting, the edge weights represent the relative importance or strength of connections between nodes, and therefore contribute directly to the identification of focal structures. By incorporating weights into the analysis, WFSA can identify influential subgroups not only by their structural position but also by the magnitude of their influence on recommendations.

In the unweighted FSA framework, the bi-level optimization treats all edges as binary connections when evaluating both group cohesion and network influence. In WFSA, these binary connections (0/1) are replaced by a weighted value (0-1) derived from the rank- and depth-weighting schemes described above. The cohesion objective is therefore computed over weighted internal connections, giving greater importance to edges that represent high-probability navigation paths. Similarly, the influence objective evaluates the weighted connectivity between the focal set and the broader network, prioritizing groups whose members are connected to the rest of the network via high-weight edges. As a result, WFSA tends to select nodes that are not merely structurally connected but are connected through pathways that users are most likely to follow, reflecting the hierarchical logic embedded in the AI recommendation system.

Jaccard Similarity

The Jaccard similarity measures the similarity between two sets by computing the ratio of the size of their intersection to the size of their union (Bag, Kumar, and Tiwari 2019). In this study, it was applied to compare the sets of nodes present in focal structures extracted by the FSA and all five WFSA methods. A higher Jaccard similarity value indicates that both methods identify largely overlapping focal structures, while a lower value indicates that the methods capture different groups of nodes. This measure allows us to assess whether FSA and the different WFSA methods identify similar or distinct focal structures within the recommendation network.

For any two focal structures A and B, the Jaccard Similarity is calculated by:

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|} \tag{9}$$

Where $|A \cap B|$ refers to the common nodes of A and B, while $|A \cup B|$ refers to the total number of unique nodes of A and B.

Structural Properties

To understand how different focal structure extraction methods influence the resulting groups, we examine key structural properties of the extracted subgraphs. These include measures such as average degree, clustering coefficient, centrality, and modularity, which describe how nodes and edges are organized.

These properties help reveal how tightly connected the identified groups are and how they are positioned within the overall network. In recommendation networks, they also provide insight into how information pathways are structured and how influence may spread. By comparing these properties across methods, we can better understand how each approach captures influential groups and their impact on overall network structure.

We examine four structural properties of the extracted subgraphs. Nodes and Edges describe the size and connectivity of the network. Average degree (2E/N for undirected graphs) captures how well-connected the typical node is. Average clustering coefficient (ranging from 0 to 1) measures the tendency of nodes to form tightly connected local neighborhoods. Average betweenness centrality captures how frequently a node lies on shortest paths between other nodes, reflecting its role as a structural intermediary. Modularity measures the strength of community structure, with higher values indicating well-defined internal clusters.

Network Resiliency Metrics

To evaluate the structural robustness of the recommendation networks, we perform targeted node removals based on the focal sets identified by FSA and the five WFSA variants. For each focal set detected by either method, all nodes in that set are removed from the network, along with their associated edges. After the removal, we measure how the network's connectivity changes using several resiliency metrics.

This process allows us to assess how removing different influential groups affects the integrity of the recommendation network. By comparing the resulting structural changes across methods, we can evaluate which focal structures play a more critical role in maintaining network connectivity. The following metrics are computed on the resulting network after the focal sets have been removed.

Isolated Nodes and Cluster Analysis: An isolated node is a node with no edges connecting it to other nodes. After removing a focal set, we count the number of nodes that become isolated. A large number of isolated nodes indicates that removing the focal set has caused many nodes to lose all their recommendation connections. In such cases, these nodes are no longer reachable from the rest of the network, suggesting a disruption in information flow.

We also examine how the network fragments into disconnected clusters after the focal sets are removed. Besides the main connected component, the number of smaller disconnected clusters is counted. Each cluster represents a group of nodes that remain internally connected but are no longer connected to the rest of the network. A larger number of such clusters indicates greater fragmentation of the network, suggesting that the removed focal set acted as a structural bridge connecting different regions of the network.

Connectivity Loss: Connectivity Loss (CL) measures the reduction in the size of the largest connected component after removing a set of nodes, normalized by the number of nodes removed. It is defined as:

$$\mathrm{CL} = \frac{|\mathrm{LCC}(G)| - |\mathrm{LCC}(G \setminus S)|}{|S|} \tag{10}$$

where $LCC(G)$ represents the size of the largest connected component in the original network, $LCC(G \setminus S)$ represents the size of the largest connected component after removing the node set S, and $|S|$ is the number of removed nodes.

Connectivity loss quantifies the average impact of each removed node on the network's overall connectivity. A higher CL value indicates that, on average, each removed node causes a larger reduction in the size of the largest connected component, suggesting greater structural importance of the removed group. This metric is particularly useful when comparing focal sets of different sizes because it measures the average impact per removed node rather than the total reduction in connectivity.

Results and Analysis

We compare focal structures extracted by unweighted FSA and WFSA across structural property comparisons and network resiliency experiments to evaluate how weighting recommendation edges by rank and depth alters structural importance and network stability.

Recommendation Graph Construction and Weight Analysis

We first constructed the recommendation networks by connecting each source–target video pair and removing duplicate edges, yielding unweighted recommendation graphs. To better capture the hierarchical nature of YouTube recommendations, we applied the proposed weighting schemes to compute edge weights based on recommendation rank and depth. This process produced weighted recommendation graphs that preserve the structural connections between videos while incorporating the relative importance of recommendation links.

Table 3: Rank weighting, depth weighting, and aggregation configurations for the five WFSA variants.

Method	Rank Weighting	Depth Weighting	Aggregation
WFSA1	Exponential decay	Exponential decay	Count weighted
WFSA2	Exponential decay	Exponential decay	Rank weighted
WFSA3	Exponential decay	Inverse decay	Count weighted
WFSA4	Linear decay	Inverse decay	Count weighted
WFSA5	Linear decay	Inverse decay	Rank weighted

Since multiple weighting strategies and aggregation schemes were introduced, we examined the resulting weight distributions to understand how the weighting approaches behave in practice. Histograms were generated for each weighting method across all datasets to observe the distribution of edge weights. During this analysis, we found that two approaches: (1) exponential rank decay with inverse depth decay using rank-weighted aggregation and (2) linear rank decay with inverse depth decay using the same aggregation rule, produced nearly identical weight distributions across all datasets. Because both methods behaved almost identically, the exponential rank-decay-inverse-depth-decay-rank-weighted approach was removed from subsequent experiments. The remaining weighting schemes were retained and labeled WFSA variants for subsequent comparative analysis, as summarized in Table 3.

Comparison of Focal Structures

After extracting focal structures using both FSA and the WFSA variants, we examined whether the two approaches identify similar sets of nodes within the recommendation networks. To evaluate this, we computed the Jaccard similarity between the top five focal structures detected by FSA and those identified by each WFSA method across all 15 datasets.

The resulting similarity scores show a strong concentration near zero, indicating that FSA and the WFSA variants typically identify different focal structures within the same network. As illustrated in Figure 2, only a small fraction of comparisons fall within the higher similarity range, suggesting that exact or near-identical overlaps between the focal structures extracted by the two approaches are relatively rare. This observation indicates that incorporating edge weights substantially changes which nodes are identified as focal structures. This result suggests that the structural influence of content in recommendation networks cannot be fully understood without considering the hierarchical signals embedded in algorithmic ranking. While FSA identifies influential groups based solely on structural connectivity, WFSA additionally considers the strength of recommendation relationships derived from rank and depth information. As a result, WFSA tends to highlight different sets of nodes that better capture the hierarchical influence patterns embedded within recommendation systems.

In the context of algorithmically curated platforms such as YouTube, these findings suggest that recommendation ranking signals can reshape which groups of content occupy structurally influential positions within the recommendation network. Consequently, weighted analysis provides additional insight into how recommendation systems may concentrate influence along particular navigation pathways in AI-driven social media environments.

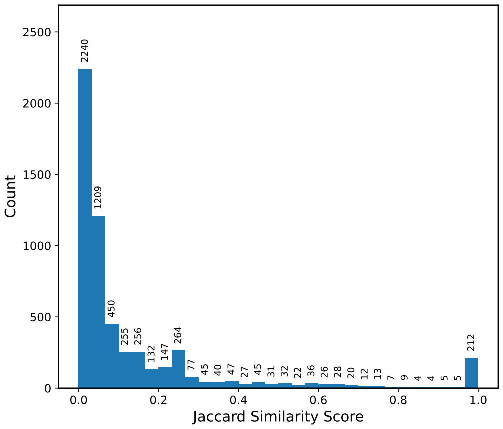

Figure 2: Distribution of Jaccard similarity scores between top-5 focal structures identified by FSA and five WFSA variants across all 15 datasets.

Structural Properties of Focal Structures

To understand how the focal structures extracted by FSA and the WFSA variants differ, we examine several structural properties of the resulting subgraphs. These include the number of nodes and edges, average degree, betweenness centrality, clustering coefficient, and modularity.

First, we look at the size of the focal structures. Figures 3(a) and 3(b) show that focal structures identified by FSA generally contain more nodes and edges than those extracted by the WFSA variants. This means that FSA tends to select larger groups of nodes within the recommendation network.

However, being larger does not necessarily mean that the nodes are strongly connected. When we examine the connectivity measures, a different pattern appears. Figures 3(c) and 3(d) show that nodes within WFSA focal structures have higher average degree and higher betweenness centrality compared to those in FSA structures. This indicates that the nodes selected by WFSA are more densely connected and occupy more important positions in the network's connectivity paths.

The difference in cohesion is also evident in the clustering

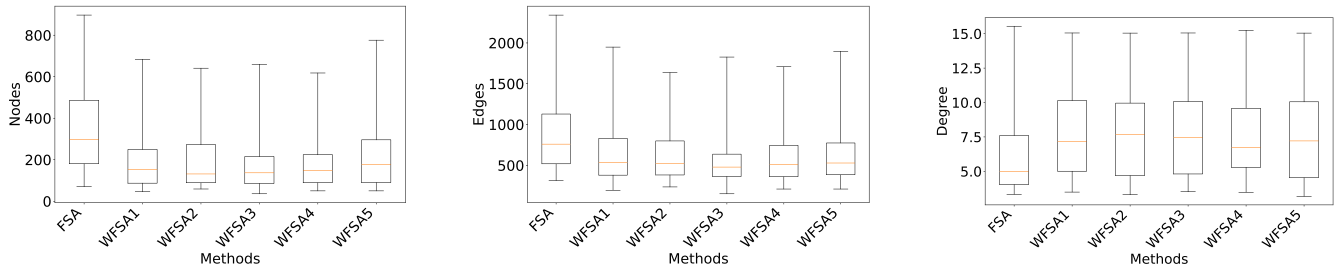

(a) Nodes

(b) Edges

(c) Average Degree

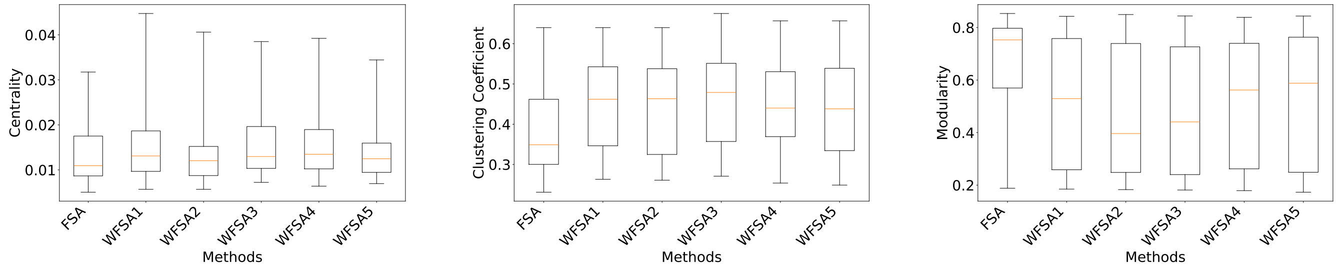

(d) Average Centrality

(e) Average Clustering Coefficient

(f) Modularity

Figure 3: Comparison of nodes, edges, average degree, betweenness centrality, clustering coefficient, and modularity of focal sets extracted by FSA and five WFSA variants across all 15 datasets.

coefficient. Figure 3(e) shows that WFSA focal structures consistently have higher clustering coefficients. This means that the nodes within these groups tend to form tightly connected neighborhoods in which many nodes are linked to one another.

We also observe a difference in modularity. As shown in Figure 3(f), WFSA focal structures exhibit lower modularity compared to those identified by FSA, indicating that WFSA captures more cohesive and internally unified groups with fewer distinct sub-communities. In contrast, the higher modularity observed in FSA suggests that its focal structures tend to include multiple internally distinct sub-communities within a single larger group. This highlights that WFSA identifies more tightly connected and homogeneous regions, whereas FSA captures broader but structurally heterogeneous sets.

The structural comparison in Figure 3 shows that WFSA identifies focal structures that are smaller but more tightly connected than those detected by FSA. While FSA extracts larger sets of nodes, the focal structures detected by WFSA contain fewer nodes but exhibit stronger internal connectivity and higher structural importance. This indicates that incorporating rank and depth information into the recommendation graph helps reveal influential groups that are not visible in the unweighted network, directly answering **RQ1** and highlighting how recommendation hierarchy shapes structural influence in algorithmically curated systems.

In recommendation platforms such as YouTube, these tightly connected focal structures can represent groups of videos that reinforce recommendation paths. Because the videos inside these structures are highly interconnected, users who follow recommendations may remain within the same part of the recommendation network. As a result, these structures can guide users toward a limited region of the recommendation graph.

Network Resiliency Assessment

To understand how important the detected focal structures are for maintaining the recommendation network, we perform a network resiliency analysis. In this analysis, the focal sets identified by each method are removed from the network, and the resulting structural changes are measured using several resiliency metrics. These include the number of nodes removed, as well as normalized measures of structural impact, such as the number of edges removed, the number of isolated nodes formed, the number of clusters formed, and the connectivity loss per removed node.

Figure 4 shows how the recommendation networks respond when focal structures detected by FSA and the WFSA variants are removed. A consistent pattern appears across all datasets. We first examine the number of nodes removed. Figure 4(a) shows that FSA generally removes a larger number of nodes compared to the WFSA variants. This means that FSA tends to identify larger focal sets within the network. However, removing more nodes does not necessarily lead to stronger disruption.

When we look at the structural impact of the removals, a different pattern becomes visible. Figures 4(b), 4(c), and 4(d) show that removing focal structures identified by WFSA leads to greater structural changes in the network. In

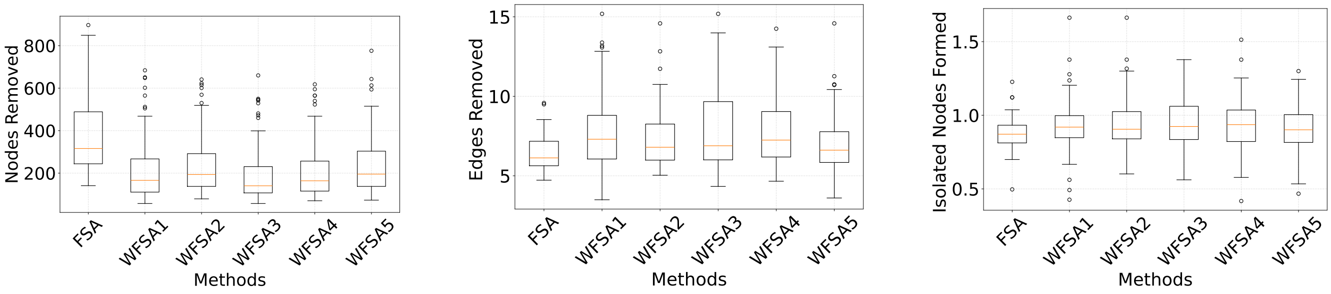

(a) Nodes removed (b) Edges removed per node (c) Isolated nodes formed per node

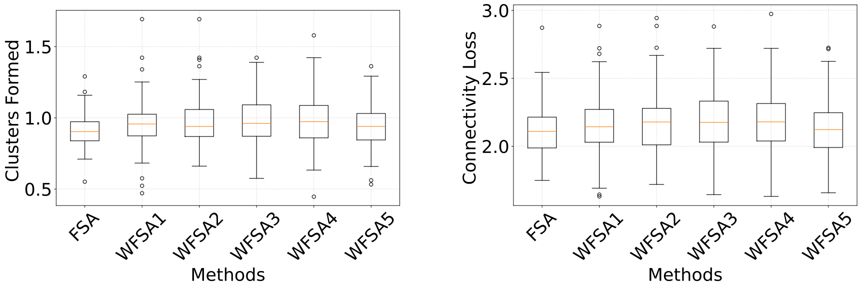

(d) Clusters formed per node (e) Connectivity loss

Figure 4: Comparison of five network resiliency metrics after targeted removal of focal sets identified by FSA and five WFSA variants across all 15 datasets.

particular, the removal of nodes identified by WFSA results in more edges removed per node, along with the formation of more isolated nodes and disconnected clusters. This indicates that each removed node has a stronger structural impact on the network, suggesting that WFSA selects nodes that occupy more critical positions within the recommendation structure.

The same trend is evident in the connectivity loss results. As shown in Figure 4(e), removing WFSA-derived focal structures results in higher connectivity loss per removed node compared to FSA. This means that the network becomes more fragmented when these nodes are removed. The relatively consistent distribution of the WFSA results across datasets also suggests that this pattern is stable and not caused by a few extreme cases.

Across all metrics, WFSA focal structures play a more important structural role in maintaining the recommendation network than those identified by FSA. Even though WFSA removes fewer nodes, the network experiences stronger disruption when those nodes are removed. This suggests that influence within recommendation systems may be concentrated within small but structurally critical groups of content that help maintain recommendation pathways. This indicates that WFSA is more effective at identifying structurally influential groups within the recommendation network, directly addressing **RQ2**.

After establishing that WFSA focal structures lead to stronger network disruption than those identified by FSA, we further compare the different WFSA weighting schemes to assess their relative effectiveness. Figure 4 reveals a clear hierarchy among the WFSA variants. Across the resiliency metrics, removing focal structures identified by WFSA3 and WFSA4 yields the most significant structural changes. In particular, Figure 4(e) shows that these methods result in higher connectivity loss per removed node compared to other variants. Between them, WFSA4 exhibits slightly stronger effects, leading to greater fragmentation. As shown in Figures 4(c) and 4(d), removing WFSA4 focal structures results in more isolated nodes and disconnected clusters, indicating that these nodes play critical roles in maintaining connections across different parts of the network. WFSA2 follows closely behind, also producing substantial structural impact.

In contrast, WFSA1 and WFSA5 lead to comparatively smaller changes in connectivity and fragmentation, suggesting that these weighting schemes are less effective at identifying structurally important nodes within the recommendation network.

These results indicate that weighting schemes combining rank importance with inverse-depth decay are more effective at identifying structurally influential focal structures. Among the evaluated approaches, WFSA4 consistently causes the strongest network disruption when its focal structures are removed, followed by WFSA3, WFSA2, WFSA1, and WFSA5. This indicates that WFSA4 most effectively captures structurally critical groups within the recommendation network, thereby answering **RQ3**.

Implications

These findings offer insights into how influence emerges in algorithmically curated systems. The hierarchical signals embedded in YouTube's recommendation system, i.e., rank and depth, change which content groups are identified as structurally influential. Unlike neutral suggestion mechanisms, recommendation algorithms rank, prioritize, and organize content in ways that shape user navigation, and these choices are reflected in network structure.

Our results show that incorporating a recommendation hierarchy reveals compact, tightly connected groups occupying disproportionately important positions in the network. When removed, the network fragments more dramatically than when FSA-identified groups are removed. In practical terms, a small set of structurally important videos may act as bridges connecting different regions of the recommendation graph, repeatedly routing users through the same content, regardless of their intent, not because the content matches their preferences, but because the network structure and ranking logic make these pathways most accessible.

For platform designers, these findings suggest that ranking functions, depth limits, and aggregation strategies have structural consequences beyond technical implementation; they shape how influence is distributed across recommendation networks.

Conclusion and Future Work

This study applies WFSA to YouTube recommendation networks, demonstrating how rank- and depth-based edge weighting captures structural influence in AI-driven recommendation systems more effectively than unweighted FSA. By incorporating edge-weighting schemes based on recommendation rank and depth, which represent key hierarchical signals guiding user navigation, we construct weighted recommendation graphs and extract focal structures that better reflect the structural realities of algorithmically curated platforms. These findings highlight the importance of incorporating an AI-driven recommendation structure when analyzing influence and information exposure in modern online platforms.

The experimental results show that incorporating recommendation hierarchy changes how influential groups are detected and is essential for accurately capturing how AI systems shape structural influence and user navigation pathways. While FSA tends to identify larger focal structures, WFSA consistently detects smaller but more tightly connected groups of nodes. Structural property analysis revealed that WFSA-derived focal structures exhibit higher connectivity and cohesion, indicating that they occupy more structurally important positions within the recommendation graph. Network resiliency experiments further confirm that removing WFSA-identified structures causes greater fragmentation and connectivity loss than removing those detected by FSA. This result indicates that WFSA more effectively captures groups that play a critical role in maintaining connectivity in the recommendation network. Among the evaluated weighting schemes, WFSA4, which combines linear rank decay with inverse depth decay and count-weighted aggregation, consistently produced the strongest structural disruption when its focal structures were removed.

These findings collectively suggest that the way an AI recommendation system organizes and presents content is not just a surface feature of user experience. It influences which groups of content become structurally prominent within the pathways available to users. Understanding the relationship between algorithmic hierarchy and structural power within recommendation networks is essential for studying the risks, rewards, and responsibilities that emerge as AI systems take on an increasingly central role in mediating human access to information.

Although this study demonstrates the advantages of incorporating recommendation hierarchy into focal structure analysis, it also has some limitations. The weighting schemes proposed here rely on recommendation rank and depth, capturing important but partial aspects of the AI-driven recommendation system's influence on human navigation and exposure patterns. Future work can extend this approach by incorporating additional behavioral signals, where available, such as watch time, user interaction rates, or click-through patterns, which may provide a more comprehensive representation of how users traverse AI-curated recommendation pathways. In addition, this study focuses on YouTube within two specific discourse domains. Applying WFSA to other recommendation platforms and broader topical contexts, including entertainment, health, or scientific information, would help assess the robustness of these findings across different AI-mediated environments. Additionally, future work could examine whether WFSA-identified focal structures exhibit greater thematic consistency than those identified by FSA, extending the content trap analysis of prior work to weighted recommendation networks. This can further improve our understanding of influence dynamics and structural vulnerabilities in algorithmically driven recommendation systems.

Acknowledgments

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

Alassad, M.; Hussain, M. N.; and Agarwal, N. 2019. Finding Fake News Key Spreaders in Complex Social Networks by Using Bi-Level Decomposition Optimization Method. In Agarwal, N.; Sakalauskas, L.; and Weber, G.-W., eds., *Modeling and Simulation of Social-Behavioral Phenomena in Creative Societies*, 41–54. Cham: Springer International Publishing. ISBN 978-3-030-29862-3.

Bag, S.; Kumar, S. K.; and Tiwari, M. K. 2019. An efficient recommendation generation using relevant Jaccard similarity. *Information Sciences*, 483: 53–64.

Bakshy, E.; Messing, S.; and Adamic, L. A. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348(6239): 1130–1132.

Bedi, P.; and Sharma, C. 2016. Community Detection in Social Networks. *WIREs Data Mining and Knowledge Discovery*, 6(3): 115–135.

Bhuiyan, M. M. I.; and Agarwal, N. 2025. Identification and characterization of content traps in YouTube recommendation network. In *The Seventeenth International Conference on Information, Process, and Knowledge Management (eKNOW)*.

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

Boak, J.; and Sanchez, F. 2025. Trump plans tariffs on Mexico and Canada for Tuesday, while doubling existing 10% tariffs on China. AP News article. Accessed: November 17, 2025.

Brin, S.; and Page, L. 1998. The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems*, 30(1): 107–117. Proceedings of the Seventh International World Wide Web Conference.

Cakmak, M. C.; Agarwal, N.; and Oni, R. 2024. The Bias Beneath: Analyzing Drift in YouTube's Algorithmic Recommendations. *Social Network Analysis and Mining*, 14(1): 171.

Covington, P.; Adams, J.; and Sargin, E. 2016. Deep Neural Networks for YouTube Recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems*, RecSys '16, 191–198. New York, NY, USA: Association for Computing Machinery. ISBN 9781450340359.

Craswell, N.; Zoeter, O.; Taylor, M.; and Ramsey, B. 2008. An experimental comparison of click position-bias models. In *Proceedings of the 2008 International Conference on Web Search and Data Mining*, WSDM '08, 87–94. New York, NY, USA: Association for Computing Machinery. ISBN 9781595939272.

Falade, T. C.; and Agarwal, N. 2025a. Optimizing Focal Toxic Structure Selection for Social Network Disruption: A WFSA-Integer Programming Approach. In *Complex Networks & Their Applications XIV: Proceedings of the Fourteenth International Conference on Complex Networks and Their Applications: COMPLEX NETWORKS 2025, Volume 1*, 133. Springer Nature.

Falade, T. C.; and Agarwal, N. 2025b. Weighted Focal Structure Analysis for Coordinated Toxicity Propagation. In *Complex Networks & Their Applications XIV: Proceedings of the Fourteenth International Conference on Complex Networks and Their Applications: COMPLEX NETWORKS 2025, Volume 1*, 120. Springer Nature.

Falade, T. C.; and Agarwal, N. 2026. Modeling Toxicity Propagation in Social Networks with Weighted Focal Structure Analysis and Monte Carlo Epidemic Models. In *Social Networks Analysis and Mining*, 297–305. Cham: Springer Nature Switzerland. ISBN 978-3-032-13821-7.

Freeman, L. C. 1978. Centrality in Social Networks: Conceptual Clarification. *Social Networks*, 1(3): 215–239.

Gan, N. 2024. This 2024 presidential election could change the world - and it's not happening in the US. CNN World News. Accessed: 9 March 2026.

Institute for Strategic Dialogue. 2024. Pulling Back the Curtain: An Exploration of YouTube's Recommendation Algorithm. https://www.isdglobal.org/publication/pulling-back-the-curtain-an-exploration-of-youtubes-recommendation-algorithm/. Accessed: 2026-03-05.

Joachims, T.; Granka, L.; Pan, B.; Hembrooke, H.; Radlinski, F.; and Gay, G. 2007. Evaluating the accuracy of implicit feedback from clicks and query reformulations in Web search. 25(2): 7–es.

Kleinberg, J. M. 1999. Authoritative sources in a hyperlinked environment. *J. ACM*, 46(5): 604–632.

Lawder, D.; and Shalal, A. 2025. Trump says he is discussing 10% tariff on China on Feb. 1. Reuters news article. Accessed: November 17, 2025.

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

Ribeiro, M. H.; Ottoni, R.; West, R.; Almeida, V. A. F.; and Meira, W. 2020. Auditing radicalization pathways on YouTube. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, FAT* '20, 131–141. New York, NY, USA: Association for Computing Machinery. ISBN 9781450369367.

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

WordStream. 2026. 32 Must-Know YouTube Statistics for 2026. https://www.wordstream.com/blog/ws/2020/02/27/youtube-statistics. Accessed: 2026-03-05.

Şen, F.; Wigand, R.; Agarwal, N.; Tokdemir, S.; and Kasprzyk, R. 2016. Focal Structures Analysis: Identifying Influential Sets of Individuals in a Social Network. *Social Network Analysis and Mining*, 6(1): 17.

Paper Checklist to be included in your paper

1. For most authors...
   - (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? Yes
   - (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? Yes
   - (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? Yes
   - (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? Yes
   - (e) Did you describe the limitations of your work? Yes
   - (f) Did you discuss any potential negative societal impacts of your work? Yes
   - (g) Did you discuss any potential misuse of your work? Yes
   - (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? Yes
   - (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? Yes
2. Additionally, if your study involves hypotheses testing...
   - (a) Did you clearly state the assumptions underlying all theoretical results? NA
   - (b) Have you provided justifications for all theoretical results? NA
   - (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? NA
   - (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? NA
   - (e) Did you address potential biases or limitations in your theoretical framework? NA
   - (f) Have you related your theoretical results to the existing literature in social science? NA
   - (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? NA
3. Additionally, if you are including theoretical proofs...
   - (a) Did you state the full set of assumptions of all theoretical results? NA
   - (b) Did you include complete proofs of all theoretical results? NA
4. Additionally, if you ran machine learning experiments...
   - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? NA
   - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? NA
   - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? NA
   - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? NA
   - (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? NA
   - (f) Do you discuss what is "the cost" of misclassification and fault (in)tolerance? NA
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...
   - (a) If your work uses existing assets, did you cite the creators? Yes
   - (b) Did you mention the license of the assets? No
   - (c) Did you include any new assets in the supplemental material or as a URL? No
   - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? Yes
   - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? Yes
   - (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR (see **?**)? NA
   - (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset (see **?**)? NA
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...
   - (a) Did you include the full text of instructions given to participants and screenshots? NA
   - (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? NA
   - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? NA
   - (d) Did you discuss how data is stored, shared, and deidentified? NA