TSINGHUA SCIENCE AND TECHNOLOGY
ISSN 1007-0214 01/01 pp1–18
DOI: 10.26599/TST.2025.9010134
Volume x, Number x, xxxx xxxx

Relational Diversification-based Anti-Mapping for Concealing User Preference

Shaokui Gu, Yang Zhang, Xiangguo Sun, Shuo Yu, Xu Yuan, Qiang Duan, Fenghua Tong, and Qingyuan Gong*

Abstract: Personalized recommendation engines often rely on preference mapping, which infers user interests from historical behaviors and social graph structures. Such mapping can trigger a content homogenization loop, gradually limiting users' exposure to diverse content and resulting in information cocoons. While existing approaches focus on diversifying content recommendations, they often overlook the role of social topology in shaping information exposure. This paper seeks to examine how the structure of social connections shapes users' exposure to diverse information and to develop measurable strategies for reducing information homogenization. We find that users with heterogeneous, strategically chosen links have diverse information exposure, which can break the information cocoon. To quantify the homogeneity of the environment, we introduce a Diversity-Scale score that quantifies both the diversity and the scope of content exposure by each user. We then establish a correlation between preference features and social connection structure, independent of explicit personal preference data, and demonstrate that strategies with diverse connections significantly weaken this correlation. We validate our approach on two real-world platforms with distinct interaction patterns (Yelp and Steam), showing that constructing diverse and dynamically heterogeneous graph structures as an anti-mapping approach is a practical and lightweight mechanism to resist preference fingerprinting and its resulting content homogeneity.

Key words: Information Cocoon; Online Social Networks; Topological Structure; Social Influence; User Behavior

1 Introduction

In online social networks, users receive content through various channels, among which recommender systems[1–4] are a typical one. Recommender systems

• Shaokui Gu is with the School of Computer Science, Fudan University, Shanghai and 200438, China. E-mail: skgu24@m.fudan.edu.cn.
• Yang Zhang is with the Anuradha and Vikas Sinha Department of Data Science, University of North Texas, Denton and 76207, USA. E-mail: yang.zhang@unt.edu.
• Xiangguo Sun is with the Department of Systems Engineering and Engineering Management, Chinese University of Hong Kong, Hong Kong and 999077, China. E-mail: xgsun@se.cuhk.edu.hk.
• Shuo Yu is with the School of Computer Science and Technology, Dalian University of Technology, Dalian and 116081, China. E-mail: yushuo@dlut.edu.cn.
• Xu Yuan is with the School of Software, Dalian University of Technology, Dalian and 116620, China. E-mail: david@dlut.edu.cn.
• Qiang Duan is with the Department of Information Technology, Pennsylvania State University Abington College, Abington and 19001, USA. E-mail: qxd2@psu.edu.
• Fenghua Tong is with the Key Laboratory of Computing Power Network and Information Security, Ministry of Education, Shandong Computer Science Center (National Supercomputer Center in Jinan), Qilu University of Technology, Jinan and 250353, China; Shandong Provincial Key Laboratory of Industrial Network and Information System Security, Shandong Fundamental Research Center for Computer Science, Jinan and 250014, China. E-mail: tongfh@qlu.edu.cn.
• Qingyuan Gong is with the Research Institute of Intelligent Complex Systems, Fudan University, Shanghai and 200438, China. E-mail: gongqingyuan@fudan.edu.cn.
* To whom correspondence should be addressed.
Manuscript received: 2025-05-11; revised: 2025-08-18; accepted: 2025-08-26

exploit rich social relationships and behavioral logs to map user preferences to help users find attractive content[5, 6], but previous research shows that these systems may reduce the diversity of content over time inadvertently[7, 8]. Recommender systems are trained to provide people's preferred objects or topics, which may confine users to familiar and repeated information. This narrowing effect has raised concerns about leading users into information cocoons. The concept of information cocoon was first proposed by Cass R. Sunstein, who described it as "a communication world where we listen only to the information we choose and feel comfortable with"[9]. The related definition of information cocoon generally refers to the limited access to diverse information, which might result in a homogeneous user experience.

Once the recommender system successfully maps user preferences, it will recommend homogeneous content to users, leading to information cocoons. For users being trapped in an information cocoon, they may not realize the extent to which they are isolated from different or opposing views, and may miss critical perspectives that could have helped make more balanced decisions, exacerbating their own homogeneous behavior[10–12] and thus highlighting their preferences, making them easier to detect by recommender systems. Existing methods[13–15] that prevent users from falling into homogeneity by increasing the diversity of recommended content, some studies based on actual platforms have found that users are still at risk of falling into homogeneity[16, 17]. For example, although improving algorithm diversity may temporarily expand user exposure, the recommender system is still based on personalized push, and relying on it may form a feedback loop, and homogeneous user interactions will strengthen algorithmic bias[18, 19]. The persistence of information homogeneity suggests that algorithmic solutions alone may be insufficient to address the problem, motivating us to investigate how the structural patterns of social connections influence users' information exposure.

The example in Fig. 1 illustrates our motivation for focusing on social connections. Although users A and B have the same number of social connections, User A's heterogeneous links enable exposure to more diverse content, whereas User B's tightly clustered network confines the user to homogeneous information. Therefore, social connections significantly influence the diversity and breadth of content that users

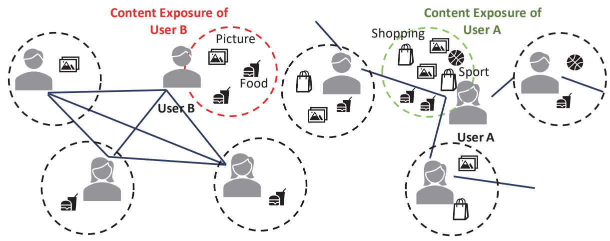

Fig. 1 shows that there is a significant difference in the diversity of the social connections of the two users. User B sits in a tightly knit ego-network whose neighbours are also densely inter-connected, forming a closed triadic core. This redundant link will deliver a small amount of homogeneous content to user B, so she is more likely to be trapped in an information cocoon. In contrast, the connections established by user A are relatively diversified. This relationship transmits a larger amount of heterogeneous information, which greatly reduces the risk of being trapped in a cocoon.

are exposed to, as well as preferences. The diversity or homogeneity of information that users access depends largely on their position in the social network. This observation motivates us to investigate how different social structures influence the formation and disruption of information cocoons and how such structures can be leveraged to promote diverse information exposure. However, understanding these mechanisms presents several technical and methodological challenges. First, existing indicators (such as entropy and normalized entropy) fail to fully balance the diversity and scale of content exposed to users, and only focus on the uniformity of content distribution. In addition, there is no relevant research that accurately represents and models the complex relationship between social connection and information cocoon.

In this paper, we focus on how social connections shape or disrupt the information cocoon effect. Our goal is to analyze the content that users are exposed to from the perspective of social connections to answer the above challenges. Specifically, (1) Existing diversity metrics fail to capture the overall exposure scale of users, limiting their ability to measure homogeneity in real settings. To address this, we propose a Diversity-Scale (DS) score that jointly considers diversity and scope in users' content exposure. Experiments show that DS effectively reveals the polarization of users' information environments and distinguishes cocooned from non-cocooned users. (2) The relationship between users' social network structure and their likelihood of falling into information cocoons remains underexplored. In particular, how to identify cocooned users purely from social topology is an open problem. To address this, we build a framework that uses network features to accurately predict cocooned users and reveal the structural indicators most associated with content homogeneity. (3) The specific structural features that shape or break information cocoons remain unclear. To further uncover these mechanisms, we analyze how key structural features interact to shape users' exposure.

2 Problem Formulation

Based on the primordial explanation of the *information cocoon* phenomenon by Sunstein[9], we define the information cocoon as a state of being confined to homogeneous information due to limited exposure to different perspectives. Formally, let $\mathcal{U} = \{u_1, \ldots, u_n\}$ represent users in a social network, and $\mathcal{C} = \{c_1, \ldots, c_m\}$ represent the universe of content items. Let $\mathcal{C}_i \subseteq \mathcal{C}$ represent the set of content items presented to u_i through its social relationships. Let $G = (\mathcal{U}, E)$ represent the social network, where E is the set of edges denoting social relationships. Let $\mathcal{F}(G, u_i) = [x_1, x_2, \ldots, x_d]$ represent structural features (e.g., Degree, PageRank, k-shell) of user u_i. To explore the information cocoon problem from the perspective of social connections, we formally propose three problem definitions:

- **P1: Measurement of Information Cocoons:** Given a set of features $\mathcal{C}_i$ for the user u_i, design a function $\mathrm{Score}(\mathcal{C}_i)$ such that u_i is in an information cocoon when

$$\mathrm{Score}(\mathcal{C}_i) \leqslant \theta,$$

where θ is a predefined threshold reflecting the minimum acceptable level of diversity. We need to define a score to accurately measure the diversity and opportunities of content that u_i is exposed to.

- **P2: Prediction of the Formation of Cocoons:** Given a $G = (\mathcal{U}, E)$, identify structural features $\mathcal{F}(G, u_i)$ (e.g., node centrality, community diversity, clustering coefficient) that are associated with low Score. Formally:

$$P(\mathrm{Score}(\mathcal{C}_i) \leqslant \theta \mid \mathcal{F}(G, u_i)).$$

- **P3: Identification of structural features that break and shape information cocoons:** Given a vector of structural attributes $\mathcal{F}(G, u_i) = [x_1, x_2, \ldots, x_d]$ (such as degree, clustering coefficient, community diversity, etc.), we want to ask how do structural features $\mathbf{x}$ influence a user's probability of being trapped in an information cocoon. We can formula this problem as:

$$P(\mathrm{Score}(\mathcal{C}_i) > \theta \mid \mathcal{F}(G, u_i)) = \sigma\left(\sum_{j=1}^{d} \beta_j \cdot x_j\right),$$

where $\sigma(\cdot)$ is a logical function, and β_j is the weight coefficient of each feature. The larger the weight value, the more significant the impact on the formation of an information cocoon.

3 Method

In this section, we propose a comprehensive framework to measure and predict information cocoons by integrating users' social structural features with content exposure characteristics, as illustrated in Fig. 2

3.1 Metric for Information Cocoons

Diversity of exposure[20] refers to the range of opinion tendencies in the content that users are exposed to on social networks, measuring whether users are exposed to diverse and heterogeneous information perspectives, which is directly related to user preferences. Ideally, users should be exposed to a wide range of opinions,

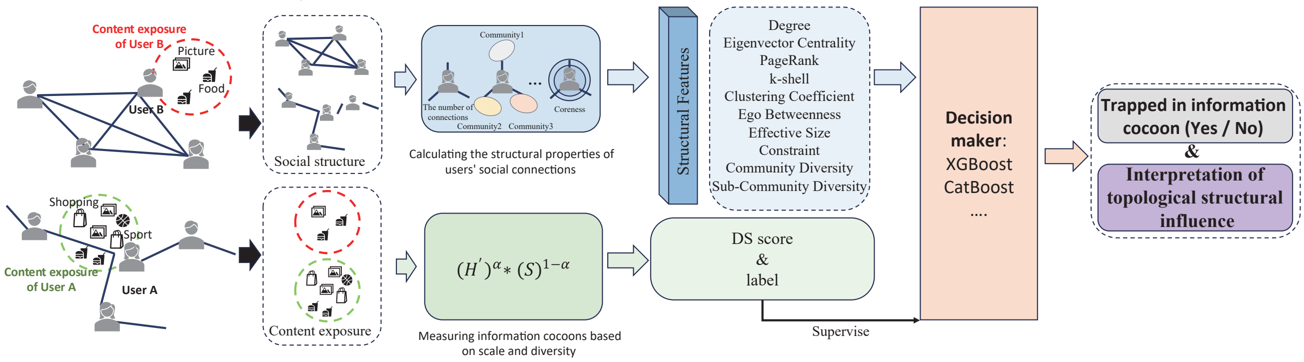

Fig. 2 The overall workflow of the proposed framework.

topics, and perspectives. However, in online social networks, this diversity is under threat, and a reduction in exposure diversity can lead to the creation of information cocoons.

One of the important factors is that when users establish relationships with other nodes, they are usually based on homogeneity in terms of preferences, opinions, interests, etc. When social connections are formed in this way, users within the group share and repeat the same opinions, reinforcing the existing preferences of group members and having little access to external information, thus forming an information cocoon.

In this state, users' content experience becomes increasingly monotonous.

To explore whether users are trapped in information cocoons, the most critical point is to explore the diversity of content they are exposed to through their social connections. Existing diversity metrics, such as entropy or normalized entropy, have clear limitations when measuring content diversity across different sample sizes. When using entropy, users exposed to a large amount of homogeneous content may have a higher entropy value[21, 22] than those with fewer, but more diverse, content experiences. This happens because, even with low diversity, entropy increases as the quantity of content exposure. On the other hand, when using normalized entropy, users with very limited but completely distinct content exposure may have a higher normalized entropy than those who interact with a large amount of diverse content (including some repetition). This is because normalization eliminates the effect of content quantity and may overemphasize the diversity of small samples.

As a result, we define users' being trapped in an information cocoon according to two complementary exposures. First, diversity measures the different content that users are exposed to. When this diversity is low, users will repeatedly see the same or highly similar content. Second, the content exposure scale measures the scale of content that users are exposed to. When the scale is low, even if the content they are exposed to is different, it will lead to limited exposure to content. Our goal is to design a comprehensive diversity metric that balances content diversity and content exposure scale, ensuring that the metric accurately reflects the extent to which users are trapped in information cocoons. To achieve this, we define two components:

- **Diversity component:** Measures the diversity of user content exposure through social connections, considering both richness and evenness. This is achieved through normalized Shannon entropy, which consists of Shannon entropy $H = -\sum_{c\in C_i} p(c)\log p(c)$ that measures the uncertainty of the content distribution and the maximum possible diversity H_{max} when all contents are equally represented:

$$H' = \frac{H}{H_{max}}, \tag{1}$$

- **Content Exposure Scale:** It represents the scale of content that users are exposed to through social relationships, and measures whether the user's access to content is limited:

$$S = \frac{1}{1 + e^{-k(\ln(N+1)-N_0)}}, \tag{2}$$

where N is the number of contents a user is exposed to, and N_0 is equal to the median of $\ln(N+1)$ of all users. We use the median instead of the mean to mitigate the impact of extreme values. k is a parameter that controls the steepness of the function curve, that is, a sensitive factor that adjusts the degree to which the user's exposure scale deviates from the baseline.

To capture both the diversity and the scale of exposure in a unified metric, we propose the following DS score:

$$\mathrm{DS} = (H')^{\alpha} * (S)^{1-\alpha}. \tag{3}$$

The DS score ranges between 0 and 1. A value approaching 0 indicates either highly homogeneous content (low H') or extremely limited exposure scale (low S), both characteristic of information cocoons. Conversely, a DS score near 1 reflects both high diversity ($H' \to 1$) and sufficient exposure ($S \to 1$). The parameter α in Eq. (3) controls the trade-off: $\alpha > 0.5$ prioritizes diversity, while $\alpha < 0.5$ emphasizes scale.

3.2 Prediction of the Formation of Cocoons

Having quantified users' information cocoon levels via the DS score in Section 3, we now explore whether structural properties of users' social connections can explain or predict their cocoon status. In this section, we first extract a set of topological features to capture the position and influence of each user in the network. Next, we introduce some classifiers to predict whether a user is trapped in an information cocoon based solely on these structural features[23].

Features. In social networks, users' information exposure characteristics can be mapped into four key dimensions through topological structure characteristics: Degree, Global Metrics, Local Metrics and Community Metrics. (1) The breadth of information exposure reflects the range of information that users can access through social links, which can be measured by Degree. (2) Global metrics include Eigenvector Centrality, PageRank, and k-shell. These metrics reflect the information aggregation ability of nodes in the global network. (3) Local metrics such as Clustering Coefficient, Ego Betweenness[24], Constraint[25, 26], and Effective Size[25, 26]. These indicators can measure the local circulation of information. (4) Community diversity can reflect the breadth and heterogeneity of users' cross-community connections:

● **Degree:** The degree of a user u_i, denoted as $\mathrm{degree}(u_i)$, represents the number of social connections and also determines the user's opportunity to access information.

● **Eigenvector Centrality:** The eigenvector centrality of a user u_i, denoted as $\mathrm{Eigenvector\ Centrality}(u_i)$, measures the influence of the user in the network by considering both direct and indirect connections, reflecting how embedded a user is within influential sub-nets that may amplify or limit exposure diversity.

● **PageRank:** The PageRank score of a user u_i, denoted as $\mathrm{PageRank}(u_i)$, is a recursive measure of the user's importance in the network based on the structure of incoming links, which correlates with the diversity of content they can encounter.

● **k-shell:** The k-shell of a user u_i is a measure of their centrality based on iterative node removal. Users with higher k-shell values are located more centrally within the core of the network, which means they may have stable access to content.

● **Clustering Coefficient:** The clustering coefficient of a user u_i, denoted as $\mathrm{CC}(u_i)$, measures the extent to which u_i's neighbors $u_k \in N(u_i)$are interconnected. High CC often means that information is constantly reinforced and echoed within a small group. On the contrary, low CC may have more cross-group connections and can bring more external content input.

$$\mathrm{CC}(u_i) = \frac{2|\{(u_j,\ u_k) : u_j,\ u_k \in N(u_i),\ (u_j, u_k) \in E\}|}{\mathrm{degree}(u_i)(\mathrm{degree}(u_i) - 1)}$$

● **Ego Betweenness:** Ego betweenness of a user u_i captures the extent to which u_i serves as a bridge or intermediary between different parts of their immediate social circle (ego network), indicating their capacity to connect distinct information clusters:

$$\mathrm{Ego}(u_i) = \sum_{s,t \in \mathrm{ego}(u_i)} \frac{\sigma_{st}(u_i)}{\sigma_{st}}$$

where σ_{st} is the number of shortest paths between s and t, and $\sigma_{st}(u_i)$ is the number of these paths that pass through u_i.

● **Effective Size:** The effective size of a user u_i quantifies the non-redundancy of their network connections, measuring how many of their neighbors $u_j \in N(u_i)$ provide unique, non-overlapping contents. Let $R(u_i) = \frac{\sum_{u_j \in N(u_i)} |\{N(u_i) \cap N(u_j)\}|}{|\{N(u_i)\}|}$ denote the average redundancy of u_i's neighborhood. Then the effective size is defined as:

$$\mathrm{Effecitve\ size}(u_i) = |\{N(u_i)\}| - R(u_i)$$

● **Constraint:** Constraint is a measure of the extent to which a user's connections are interconnected, limiting the user's ability to access diverse information.

$$\mathrm{Constraint}(u_i) = \sum_{u_j \in N(u_i)} \left(\frac{p_{ij} + \sum_{u_k \in N(u_i)} p_{ik} p_{kj}}{\mathrm{degree}(u_i)} \right)$$

● **Community Diversity and Sub-Community Diversity:** The Community Diversity measures the heterogeneity of a user's social connections across communities (via the Louvain algorithm). Let M_i denote the i-th community connected to user u, and w_i represent the weight (e.g., size) of M_i. The Community Diversity index is defined as:

$$\mathrm{Community\ Diversity}(u_i) = \sum_{i=1}^{n} \left(\frac{w_i}{W} \right)^2,$$

where $W = \sum w_i$. A high $\mathrm{Community\ Diversity}(u)$ indicates balanced connections to multiple communities, reducing cocoon risks.

The Community Diversity within finer-grained subcommunities is calculated by the Louvain algorithm to capture local homogeneity. Users with low subcommunity diversity may have access to a diverse global community, but are still trapped in a local cocoon.

Classifiers. We employ tree-based classifiers, including decision trees[27], random forests[28], and boosting methods such as XGBoost[29] and

CatBoost[30]. The classifier aims to determine whether a user is trapped in an information cocoon. Tree-based methods gradually split the feature space, so they are able to naturally find the most informative threshold for each structural signal. Since the formation of a cocoon usually depends on the synergy of multiple network features, the tree can capture these cross-feature interactions through different branches without any manual feature engineering. In addition, tree-based models provide transparent decision rules, which are crucial for explaining how structural features influence cocoon formation. These models naturally rank features by information gain, aligning with our goal of identifying critical topological factors (see Section 3.3).

3.3 Interpretation of Topological Structural Influence

Section 3.2 established that tree-based classifiers can accurately predict whether a user is trapped in an information cocoon from structural features. We now show why the model makes its decisions and how structural features shape and disrupt the information cocoon with SHAP[31]. By taking a set of features and the classifier as input, the SHAP library outputs the average SHAP value for each feature. For a trained tree-based model $f(x)$ and the feature vector $\mathcal{F}(G, u_i) = [x_1, x_2, \ldots, x_d]$ of user u_i, where $d = 10$, these 10 features are [Degree, Eigenvector Centrality, PageRank, k-shell, Clustering Coefficient, Ego Betweenness, Effective Size, Constraint, Community Diversity, and Sub-Community Diversity] introduced in section 3.2. SHAP assigns each feature x_j a contribution score $\phi(x_j)$, indicating how much x_j decreases the probability of being trapped in cocoons:

$$f(x) = \phi_0 + \sum_{j=1}^{d} \phi(x_j) \tag{4}$$

where ϕ_0 is the model's expected output $\mathbb{E}(f(\mathbf{x}))$. The values $\phi(x_j) \in \mathbb{R}$ is the contribution of representation x_j (e.g., high degree) to the prediction result. SHAP value be computed as:

$$\phi(x_j) = \sum_{D \subseteq \mathcal{F} \setminus \{j\}} \frac{|D|!(d-|D|-1)!}{d!} \left[f(D \cup \{j\}) - f(D) \right], \tag{5}$$

where D is a subset of $\mathcal{F} \setminus j$. $f(D)$ represents the model prediction value when only the feature subset S is used. Formula 5 reveals how much gain a feature (e.g., degree) brings to the model prediction, which is obtained by calculating the difference in gain when the feature is removed and added.

To further analyze the interactions between features, SHAP defines a symmetric interaction value $\phi_{j,k}$ that quantifies the joint impact of features j and k. For any two j, k, the formula for their interaction is:

$$\phi(x_j, x_k) = \sum_{D \subseteq F \setminus \{i,j\}} \frac{|D|!(d-|D|-2)!}{2(d-1)!} \Big(f(D \cup \{j, k\}) - f(D \cup \{j\}) - f(D \cup \{k\}) + f(D) \Big) \tag{6}$$

where D is the subset that does not include j and k. The terms inside the brackets represent the joint marginal contribution of features j and k, that is, their synergistic effect. Formula 6 reveals whether the two structural factors acting together increase (e.g., high degree & low community diversity) or reduce (e.g., high degree & low clustering coefficient) the risk of being trapped in an information cocoon.

4 EMPIRICAL ANALYSIS

In this section, we present our empirical analysis to reveal the relationship between users' social connections and information cocoons. Specifically, we aim to answer the following research questions:

- **RQ1:** How can we construct a unified metric that captures both the scale and diversity of information solely based on the content users are exposed to through their social connections, in order to accurately assess whether they are trapped in information cocoons?
- **RQ2:** In the absence of explicit representations of user interests or preferences, is it possible to predict the likelihood of a user being trapped in an information cocoon solely based on network structural signals?
- **RQ3:** What kind of social connection strategies can leverage structures to increase the exposure to diverse information, reduce the probability of preference mapping, and thus prevent users from falling into information cocoons?

4.1 Dataset and Parameter Setting

We use two datasets from real-world platforms, i.e., Yelp and Steam, which are summarized in Table 1.

Yelp Dataset. Yelp is a life service platform based on user reviews, providing business information, ratings and reviews to help users discover and select venues such as restaurants and shops. Friendships can be established between different users. The Yelp

Table 1 Dataset statistics. We count the number of users in the largest connected component, the total number of relations, and their density

Dataset	Users	Relations	Density
Yelp	892,152	7,298,492	1.834×10^{-5}
Steam	1,977,628	5,176,970	2.647×10^{-6}

dataset used in this paper was released in July 2023 by Yelp Inc ※. It contains information on 1,987,897 users and their 6,990,280 reviews across 150,346 businesses. We selected three subsets of this dataset for analysis: User, Business, and Review. We used the friend relationships provided by the user to construct the Yelp social network graph. This means that connected nodes are friends with each other, and the graph is undirected. We selected 308,002 users with three or more comment records and their neighbors from the largest connected component for analysis.

Steam Dataset[32]. Steam is a digital distribution platform for electronic games launched by Valve and is one of the largest comprehensive digital distribution platforms in the world. Players can purchase, download and discuss games on this platform. The raw data† is crawled by Mark et al.[33] through Steam Web API‡. It contains information on 3,908,744 users and their 4,079 games across 22 genres. We used the friend relationships provided by the user to construct the Steam social network graph and analyzed 1,977,628 users from the largest connected component who had three or more game records.

Parameter Setting. In this paper, we set the $k = 5$ and $\alpha = 0.5$ to the metric (3) for all datasets. $k = 5$ was chosen because it is sensitive enough to respond to deviations around the median under most exposure distributions, but not overly amplified by small fluctuations. $\alpha = 0.5$ means giving equal weight to both to balance diversity and exposure scale.

4.2 Discovery of the Information Cocoon (RQ1)

This section first demonstrates the effectiveness of our proposed metric in capturing content diversity and exposure opportunities. We also explore how information cocoons influence user behavior by analyzing user-neighbor homogeneity, comparing our metric with existing indicators, and examining user behavior on social platforms.

※https://www.yelp.com/dataset

†https://steam.internet.byu.edu/

‡https://developer.valvesoftware.com/wiki/Steam_Web_API

4.2.1 Content Exposure Scale

We first studied the amount of content that users were exposed to on social platforms, as well as the logarithms of these amounts, as shown in Fig. 3. We found that the distribution of the amount of content that users were exposed to on both platforms was similar, with most users having a long-tailed distribution of content exposure. In order to be consistent with the definition of the subsequent Exposure Scale metric, after taking the logarithm of these values, we can clearly see the overall trend and concentration of the data. We can see that the distribution on both platforms is close to a Gaussian distribution. This suggests that despite differences in platform design and user behavior, there is a consistent pattern across platforms, that is, most users are not exposed to much content, while only a few users have significantly higher content exposure. Then we used the metric (3) to calculate the extent to which users are trapped in information cocoons based on the content they are exposed to. As shown in Fig. 4, the density reflects the proportion of users with each given DS value. The results indicate a clear polarization in content diversity among users on both platforms. In other words, the distribution of content diversity that users are exposed to through social connections exhibits two extremes: under our metric, a large proportion of users are exposed to relatively diverse content, while another large proportion is exposed to homogeneous content. This finding is crucial because it highlights the polarized trend in user content exposure across platforms. Furthermore, it demonstrates the effectiveness of our metric in capturing these patterns. The consistent appearance of this polarization on both platforms further validates that our metric can measure the extent of information cocoons and proves its robustness across different social platforms.

4.2.2 Definition of Users Trapped in Information Cocoons

We defined the users in the bottom 20% of calculated DS values as those trapped in an information cocoon, and the top 20% as those not trapped in an information cocoon, which includes users at the extremes. Then, we compare other metrics used to measure diversity or information cocoons, e.g., Normalized Entropy[16] and Entropy[34]. First, we calculated the overlap between users trapped in information cocoons and users not trapped in information cocoons, as screened by various

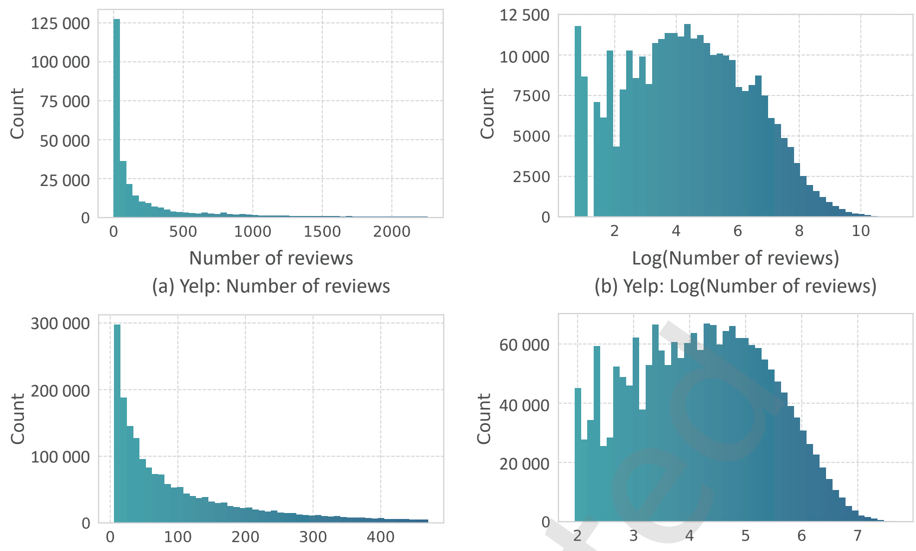

(a) Yelp: Number of reviews (b) Yelp: Log(Number of reviews)
(c) Steam: Number of games (d) Steam: Log(Number of games)

Fig. 3 Distributions of User Content Exposure on Yelp and Steam.

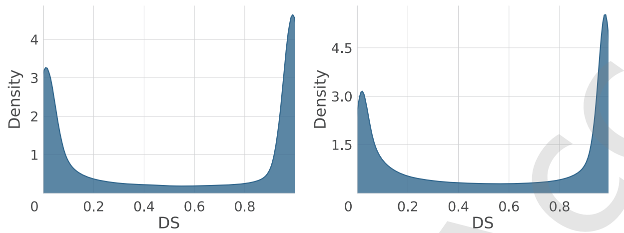

(a) Yelp: User DS value distribution (b) Steam: User DS value distribution

Fig. 4 Distribution of Information Cocoon DS Value Across Platforms.

metrics. The results are shown in Fig. 5. We observe clear differences among the users identified by the three metrics. Users selected based on entropy show a relatively high overlap with those identified by our metric across both datasets, whereas users screened using normalized entropy exhibit a much lower overlap rate.

4.2.3 Behavior Diversity of Users Trapped in Information Cocoon

To systematically explore the impact of the information cocoon on user behavior, we analyze the diversity of both users trapped in and not trapped in information cocoons using three different cocoon metrics (DS, normalized entropy, and entropy), as shown in Fig. 6. By applying these three metrics, we observe the behavioral patterns of users trapped in and not trapped in information cocoons on both the Yelp and Steam platforms.

On Yelp, regardless of which metric is used for classification, most users (both trapped in and not trapped in information cocoons) have an average rating

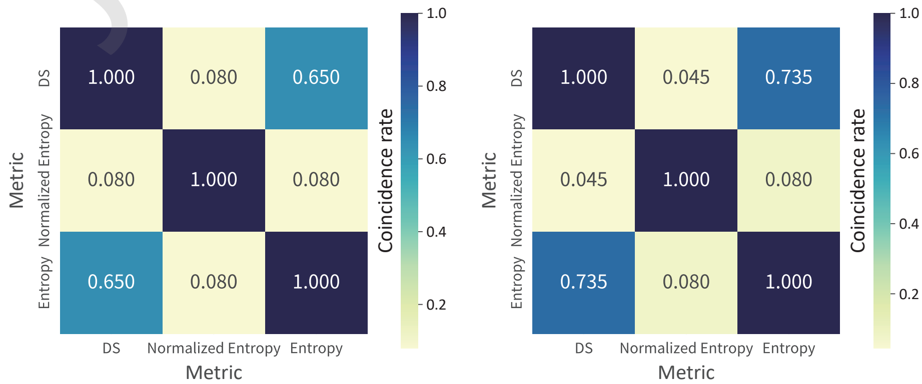

(a) Yelp (b) Steam

Fig. 5 Overlap Between Users Trapped and Not Trapped in Information Cocoons Identified by Different Metrics.

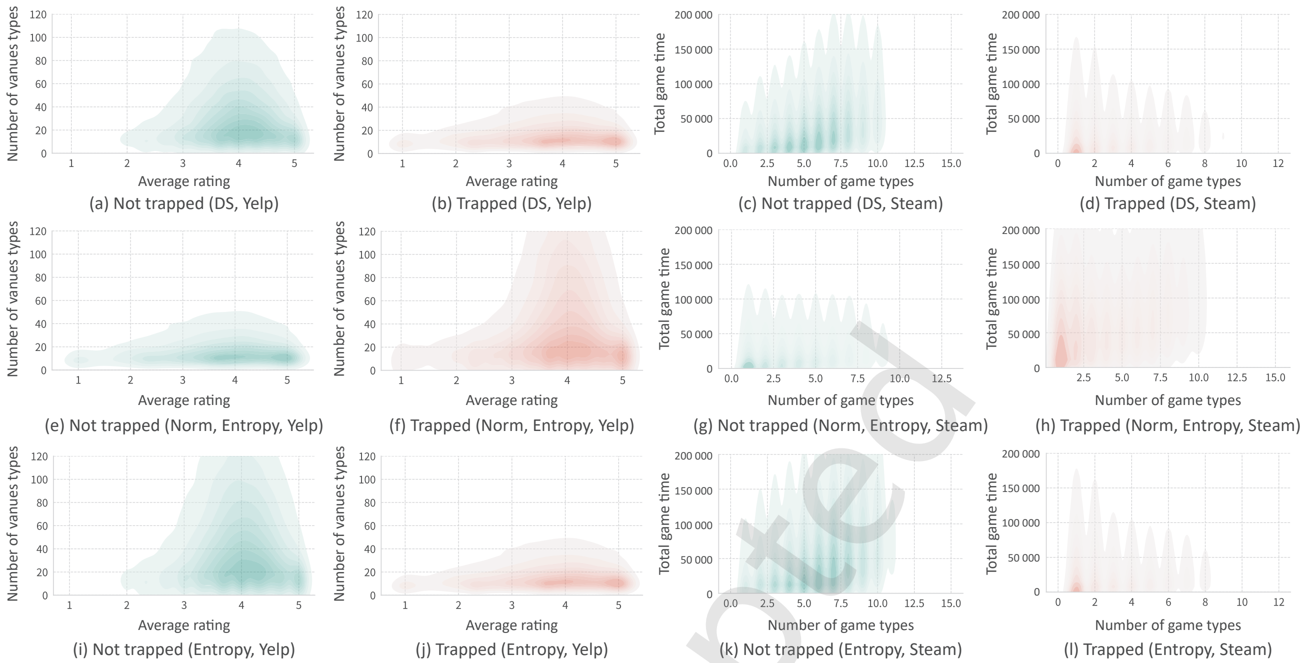

Fig. 6 Behavioral Patterns of Trapped and Not Trapped Users on Yelp and Steam Based on Different Cocoon Metrics.

concentrated around four stars. However, users identified as not trapped in information cocoons by DS and entropy consistently visit a greater variety of venues, while users trapped in information cocoons tend to restrict their activities to fewer and more homogeneous places. In contrast, users identified using normalized entropy unexpectedly exhibit the opposite pattern: users trapped in information cocoons display a higher diversity of venue visits than those not trapped in information cocoons. This is a counterintuitive and unintuitive result, suggesting that normalized entropy may not accurately capture users' true information cocoons due to its failure to account for exposure scale.

On Steam, according to both DS and entropy, users not trapped in information cocoons engage with a wider range of game types, with participation density peaking around six different types, while users trapped in information cocoons overwhelmingly focus their attention on only one or two game types. Furthermore, users trapped in information cocoons tend to accumulate longer total game time within these limited categories, indicating repetitive consumption. In contrast, users not trapped in information cocoons distribute their play time across a more diverse selection of games, suggesting greater curiosity and willingness to explore new content. Similarly, the behavioral patterns of users identified using normalized entropy are inconsistent and confusing.

4.2.4 Homogeneity Measure with Neighbors

Related studies have mentioned that the behaviors of users and their neighbors in social networks should be heterogeneous[35]. However, users trapped in information cocoons have a higher homogeneity with their environment[16]. Therefore, based on the formula in the literature[35], we propose the following metric to measure the homogeneity of users and their social connections in order to interpret their interactions with their environment and indirectly justify our metric.

Let $\mathcal{U} = \{u_1, u_2, \ldots, u_n\}$ $(|\mathcal{U}| = n)$ and $\mathcal{C} = \{c_1, c_2, \ldots, c_m\}(|\mathcal{C}| = m)$ represent the set of users and the set of items, respectively, where $|\cdot|$ denotes the number of elements in a set. For the Yelp platform, c represents the venue corresponding to the review, and for the Steam platform, c represents the game. Given a pair of friend users u_i and u_j, with their respective subsets of items $\mathcal{C}_{u_i}$ and $\mathcal{C}_{u_j}$, we define the relative homogeneity ratio as

$$rh_{(i,j)} = \frac{\left|\mathcal{C}_{u_i} \cap \mathcal{C}_{u_j}\right|}{\left|\mathcal{C}_{u_j}\right|}, \tag{7}$$

where $rh_{(i,j)} \in [0,1]$ indicates how similar user u_j's preferences are to those of user u_i. A high homogeneity ratio $rh_{(i,j)} \to 1$ suggests strong homogeneity, meaning that u_j cannot offer much diverse content to u_i. Based on this, we can calculate the relative homogeneity between u_i and all of their neighbors, and define the homogeneity ratio of the network as:

$$\mathcal{RH}_s = \frac{1}{K} \sum_{(i,j)\in\{I_{(i,j)}=1\}} rh_{(i,j)}, \tag{8}$$

where K is the number of edges (friendship pairs) in the group, and $I_{(i,j)} = 1$ indicates that u_i and u_j are friends. $\mathcal{RH}_s$ reflects the overall homogeneity of the social network.

We use the metric (8) to calculate the relative homogeneity ratio between users and their neighbors, as shown in Fig. 7. We observed that for all metrics on both platforms, users who are not trapped in information cocoons have lower experience homogeneity with their neighbors than users who are trapped in information cocoons. This suggests that users who are exposed to more diverse and heterogeneous content are less likely to be limited to homogeneous experiences than their neighbors. The homogeneity between users who are not trapped in information cocoons and their neighbors selected by our proposed metric is lower, and the homogeneity between users who are trapped in information cocoons and their neighbors is the same as or lower than that selected by other metrics. Homogeneity indicates that the content that one is exposed to is mostly from the same type of topic. The higher the homogeneity, the easier it is to fall into information cocoons, which indirectly explains the rationality of our metric. In addition, we also counted the number of venues or games that these users could access through their neighbors, as shown in Fig. 8. It can be seen that the number of contacts on the two datasets through our designed metric is moderate. This shows that the metric we designed is not affected by the number of contacts. The remaining two metrics, entropy, ignore the consideration of diversity due to the influence of quantity accumulation. Normalized entropy ignores quantitative considerations, so it shows the counterintuitive phenomenon that users trapped in information cocoons are exposed to more content information than those who are not trapped in information cocoons.

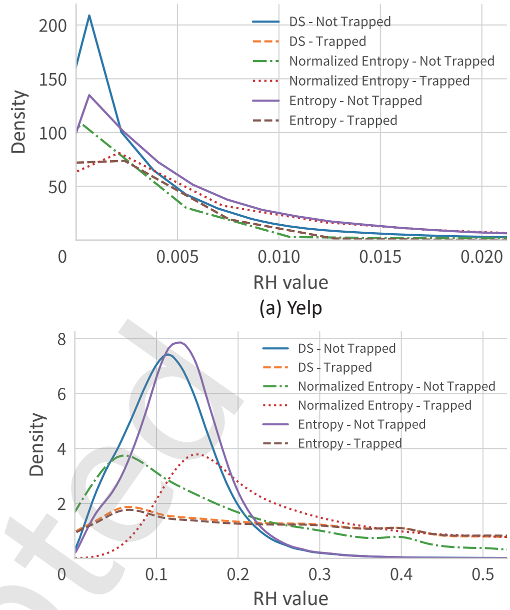

(a) Yelp

(b) Steam

Fig. 7 Density of Relative Homogeneity Ratio Between Users' Own Experiences and Their Neighbors on Yelp and Steam.

4.3 Predicting Information Cocoons from Social Structural Features (RQ2)

In this section, we focus on whether we can predict whether a user will be trapped in an information cocoon using only structural information. We trained a set of classifiers based on structure and quantified their effectiveness on the Yelp and Steam platforms.

The input of each model is the feature vector introduced in Section 3.2, covering the three key dimensions of breadth, homogeneity, and diversity of information exposure. Users are labeled as trapped in information cocoons (the lowest 20% of DS scores) or

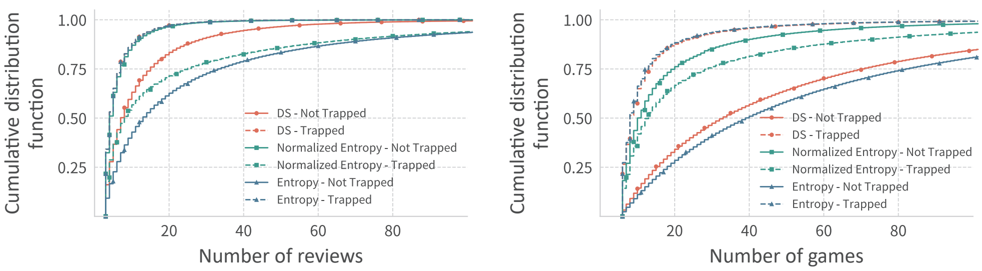

(a) Cumulative distribution of the number of reviews

(b) Cumulative distribution of the number of games

Fig. 8 User Activity Comparison Between Trapped and Not Trapped Users on Yelp and Steam.

not trapped in information cocoons (the highest 20% of DS scores). We randomly split the data (80% for training and 20% for testing) and evaluate four tree-based models: decision trees, random forests, XGBoost, and CatBoost. Table 2 reports Precision, Recall, F1-score, and AUC.

All four classifiers achieve strong discrimination, but XGBoost consistently performs best, reaching AUC 0.9991 on Yelp and 0.9644 on Steam. Even the singular DT achieves F1 values around 0.93, suggesting that simple, interpretable rules extracted from the social topology are sufficient to label most cocoons. On Steam, performance drops slightly due to its sparser social connection graph, but the model still achieves an AUC of over 0.96, demonstrating that the approach generalizes across platforms with different social dynamics and effectively predicts a user's likelihood of being trapped in an information cocoon based on structure.

Table 2 Model Comparison Results

Models	Precision	Recall	F1-score	AUC
		Yelp		
XGBoost	**0.9856**	**0.9877**	**0.9867**	**0.9991**
CatBoost	0.9844	0.9865	0.9854	0.9990
Decision Tree	0.9782	0.8940	0.9342	0.9762
Random Forest	0.9837	0.9132	0.9471	0.9901
		Steam		
XGBoost	**0.9227**	**0.8812**	**0.9015**	**0.9644**
CatBoost	0.9195	0.8769	0.8977	0.9618
Decision Tree	0.8273	0.7980	0.8124	0.8945
Random Forest	0.8215	0.7948	0.8079	0.8955

4.4 Social Connections Shape and Disrupt Information Cocoons (RQ3)

In this section, we analyze which structural features are most influential in shaping or disrupting information cocoons. We use the SHAP mentioned in section 3.3 to interpret the output results of the classifier in section 4.3. The results for Yelp and Steam are shown in Fig. 9. From the results, we observe that four features, Eigenvector Centrality, Degree, Community Diversity, and Sub-Community Diversity consistently emerge as important across both datasets. Notably, when users have high Eigenvector Centrality and Degree values, these features tend to have a significant impact on the classifier. This leads to a straightforward but critical conclusion: the number of a user's connections and the

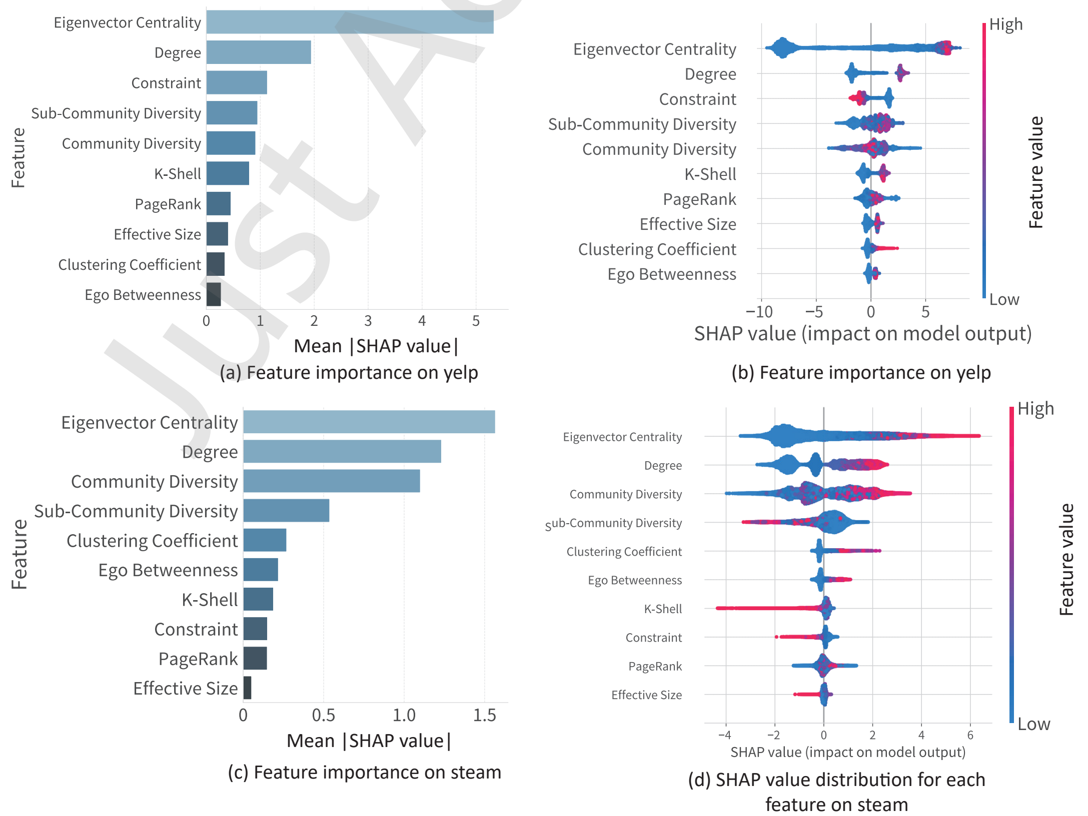

Fig. 9 SHAP Value Analysis of Features Impacting Information Cocoon Likelihood on Yelp and Steam.

importance of their neighbors are key factors influencing whether a user becomes trapped in an information cocoon.

Additionally, our analysis reveals contrasting effects for Community Diversity and Sub-Community Diversity. A higher Community Diversity value, meaning a user's connections are distributed across many different communities, significantly reduces the likelihood that the user will fall into an information cocoon. This finding suggests that bridging across diverse communities allows users to access a wider range of information sources, mitigating the risk of content homogeneity and supporting information diversity.

On the other hand, a higher Sub-Community Diversity value, which measures a user's diversity of connections within sub-communities defined by a finer-level community partition, actually increases the probability of becoming trapped in an information cocoon. This somewhat counterintuitive result may arise because users who concentrate their connections within various subgroups of a single larger community, even if these subgroups appear diverse at the micro level, are still confined to a relatively homogeneous macro environment. In other words, having high diversity within just one community can still result in information cocoon effects, since users miss out on broader perspectives that would be gained by connecting across entirely different communities.

Surprisingly, we found that users with a large number of social connections may still have an increased probability of falling into information cocoons, while users with fewer connections may avoid becoming trapped in information cocoons. To further investigate this phenomenon, we analyzed how the interactions between features affect the classifier's decision, as shown in Fig. 10. Specifically, we observed that when the degree is within a small range, the increase in the number of social connections and the clustering coefficient can turn from a negative impact on the classifier's prediction to a positive impact. Generally, a high clustering coefficient means that the user's friends are also closely connected to each other, which leads to information redundancy because the diverse information received by the user may come from the same small group, thus shaping a cocoon. However, as the number of connections and the clustering coefficient further increase, users begin to connect to different communities or groups, which breaks the original closed loop and brings more diverse content exposure to disrupt the cocoon.

In addition, as the number of social connections increases, community diversity becomes a key factor in determining whether a user will fall into an information cocoon. We found that for users with higher degrees, users with higher community diversity always contribute positively to the model's predictions. On the other hand, users with lower community diversity show the opposite effect. Each group has its own unique information and opinions, and the increase in the number of social connections and the number of connected groups greatly increases the opportunity for users to be exposed to diverse content. This is because community diversity essentially reflects the connections of users across different groups. In contrast, when a user's community diversity is low, the user's connections are more concentrated in a single group, and the information sources in these groups are often repetitive and homogeneous. Therefore, even though the number of social connections of a user may be high, such a structure may still shape an information cocoon due to the lack of cross-community connections.

Finally, we further investigate the structural properties of user neighborhoods to understand what types of neighbor structures reduce the likelihood of users falling into information cocoons, as shown in Fig. 11. Specifically, for each user, we calculate the average values of four key structural features among their neighbors: PageRank, ego betweenness, degree,

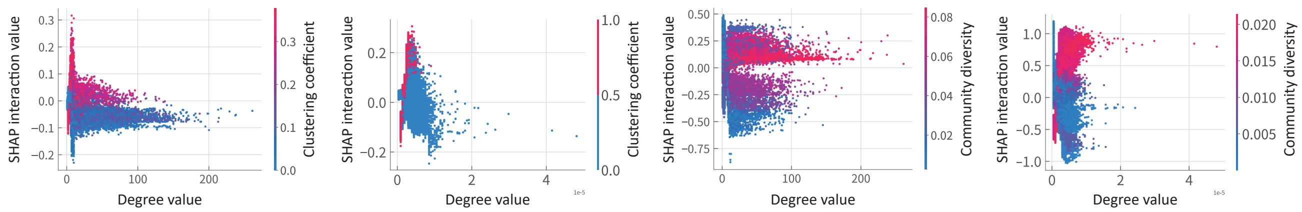

(a) Degree vs. clustering coefficient (Yelp) (b) Degree vs. clustering coefficient (Steam) (c) Degree vs. community diversity (Yelp) (d) Degree vs. community diversity (Steam)

Fig. 10 Interaction Effects of Degree and Network Diversity Measures on Yelp and Steam.

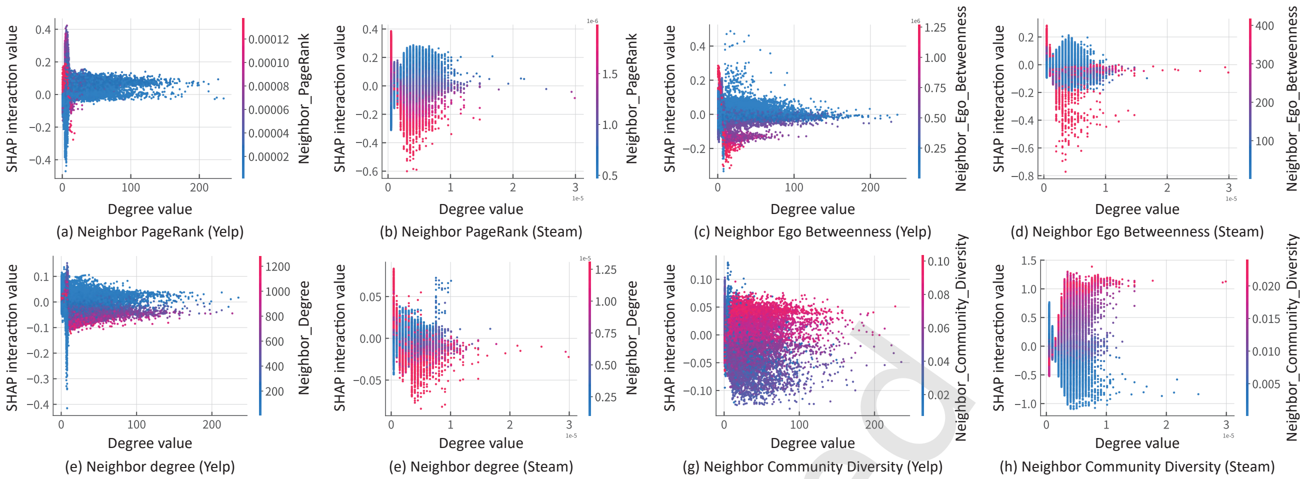

Fig. 11 Interaction Effects Between User Degree and Neighbor Structural Features on Yelp and Steam.

and community diversity. These metrics respectively represent the influence and centrality of neighbors (PageRank), their ability to act as local bridges (ego betweenness), the overall social connectivity of neighbors (degree), and the diversity of communities that neighbors are connected to (community diversity). We then examine how the interactions between these average neighbor features and a user's own degree influence the classifier's prediction.

Our analysis reveals several important patterns. When a user has a small degree (i.e., relatively few social connections), having neighbors with higher average PageRank, ego betweenness, or degree tends to decrease the probability of the user falling into an information cocoon. This suggests that when users themselves are less connected, the presence of influential, well-connected, and strategically positioned neighbors helps to provide access to more diverse information sources, reducing content homogeneity and the risk of cocooning.

However, as a user's degree increases, the situation changes. For users with a higher number of social connections, having neighbors with high average PageRank, ego betweenness, or degree actually increases the likelihood of being trapped in an information cocoon. This may be because, as users become highly connected, their network can become densely clustered and overlapping, leading to information redundancy and reinforcing the cocoon effect, especially if their many connections are within similar groups.

In contrast, as user degree increases, higher average community diversity among neighbors continues to lower the probability of falling into an information cocoon. This indicates that, regardless of how many connections a user has, the key to breaking the cocoon lies in whether their neighbors are themselves connected to a wide range of diverse communities. When neighbors span across many different communities, they serve as bridges to novel and heterogeneous content, which continually introduces the user to fresh information and perspectives.

Based on the above research results, we can offer some practical suggestions to help users avoid information cocoons and conceal preferences:

Observation 1: How to break the information cocoon for new users? For new users, a practical strategy is to recommend connections with high PageRank, ego betweenness, or degree, as these are influential and well-connected individuals who are socially close but not yet directly connected. Connecting with such users can expose newcomers to a wider range of content and communities, helping them build a diverse social network from the start.

Observation 2: How to break the information cocoon as social connections continue to increase? For users whose number of social connections is gradually increasing, it is crucial to monitor not only the growth in the number of connections but also the diversity of the communities to which these connections belong. If a user's connections are increasingly concentrated within a single community or among neighbors with low community diversity, the risk of falling into an information cocoon rises. In these cases, the platform should actively recommend users who are embedded in other distinct communities, thereby encouraging cross-community connections. This approach helps broaden users' exposure to new information sources and perspectives, ultimately preventing the formation of information cocoons as their networks expand.

5 Related Work

5.1 Information Consumption and Information Cocoons

There are several phenomena related to information cocoons in social networks that hinder people's exposure to diverse content. Phenomena closely related to information cocoon include echo chambers[36] and filter bubbles[37], but the mechanisms of each phenomenon are different. Information cocoon refers more broadly to the low diversity of information exposure and the narrow information perspective. *Filter bubbles*[37] occur when recommendation algorithms tailor content based on users' past behavior, gradually excluding dissimilar content and reinforcing users' existing preferences. *Echo chambers*[36] occur when social connections converge (i.e., users primarily connect with like-minded people) and information spreads within closed communities, reinforcing existing beliefs or ideologies through repeated interactions. These phenomena may be caused by a variety of reasons, such as recommender systems or social connections, which will affect users' information consumption. Many existing studies have shown that user behavior is influenced by social connections[30, 38−40].

Many existing works have researched the information cocoon and similar phenomena[9, 16, 38, 41−43]. For example, a study[17] on the Taobao e-commerce platform explored the relationship between the frequency of user clicks on items in the recommender system list and the diversity of the system's content. The study showed that a tendency towards echo chambers exists in personalized recommender systems, especially when users exhibit purchasing behavior, which intensifies this tendency. Another study[44] on the Spotify dataset demonstrated that songs recommended by algorithms based on relevance are very effective in meeting users' needs in the short term. However, this recommendation algorithm is associated with a reduction in the diversity of songs. A study[45] indicates that the community structure of networks becomes stronger over time, increasing the separation between communities. In cases where communities indeed align with viewpoints, this separation can be interpreted as polarization.

5.2 Identifying and Breaking the Information Cocoon

Despite the emergence of complex recommendation technologies and the increasingly diverse channels for users to obtain information, many users are still limited to a narrow range of content. With the rise of awareness about the dangers of information cocoons and similar phenomena. Some works have measured information cocoons and designed a series of methods to prevent users from falling into homogeneous information environments and promote users to discover diverse content, aiming to encourage users to actively seek diverse perspectives.

Identifying the information cocoon of an individual in a social network mainly requires a reasonable measurement of the homogeneity of its environment. The work[16] quantified the homogeneity of user video consumption behavior over time by coverage (calculating the number of different categories of user interaction) and entropy (measuring the uniformity of the distribution between these categories), and measured the degree of information cocoons on the short video platform (Kuaishou). The study considered the quality and diversity of content users are exposed to separately, but does not combine the two factors. The work[34] calculated the diversity of the information distribution received by each individual from its neighbors based on the Shannon information entropy formula, and then expressed the overall information diversity as the average of the information entropy of individuals in the entire network, but it does not take into account the scale of information that users are exposed to.

In terms of breaking the information cocoon, the SERAL[13] framework uses aligned large language models to inject serendipity into the industrial recommendation process, significantly improving the exposure of new products. GS2-RS[14] uses twin generative adversarial networks to create virtual neighbors, enrich users' fine-grained preference profiles, and alleviate cold start and filter bubble effects. CIRS[15] combines causal user modeling with offline reinforcement learning to dynamically counteract the feedback loop of overexposure to familiar content. In addition to recommendation algorithms, qualitative studies have also found contextual factors, user characteristics, and motivations that cause individuals to actively seek contradictory information to break their own cocoons[46]. At the macro level, a four-dimensional intervention path covering information ontology, technical governance, algorithmic fairness, and social cohesion has been proposed to dismantle information cocoons in emergency situations[47]. HearHere[48] is an AI-driven interface that quantifies and visualizes news bias, supports reflective exploration, and significantly mitigates the echo chamber effect in news consumption.

6 Discussion and Conclusion

Recommender systems map user preferences through their social relationships, thereby recommending homogeneous content to users, causing them to fall into an information cocoon. In this study, we successfully identified whether users on Yelp and Steam platforms were trapped in information cocoons by leveraging topological features. Specifically, we defined a new metric for measuring information cocoons, which allowed us to classify users into those trapped in information cocoons and those not. Through this analysis, we found a clear polarization in content exposure, with users being exposed to highly diverse content and those being restricted to highly homogenized content.

By calculating various structural indicators and forming feature vectors, we trained a classifier to predict whether a user was trapped in an information cocoon. We found that the likelihood of being trapped in an information cocoon is closely related to the diversity of the connected group, the number of neighbors, and their importance. In addition, we explored the relationship between information cocoons and user behavior and found that users trapped in information cocoons had a higher relative homogeneity ratio with their neighbors and were significantly less active and engaged on the platform. This finding suggests that homogenized content exposure has the potential to have a negative impact on user engagement and platform growth.

This study has important implications for the design of platform recommendation algorithms and user engagement strategies. By identifying and solving the information cocoon problem, platforms can promote more diverse content exposure to each other by controlling the overall structure of social networks, thereby increasing user engagement and the long-term growth of the platform.

Acknowledgment

This work has been sponsored by National Key R&D Program of China under Grant (No.2022YFB3102901), National Natural Science Foundation of China (No.62102094), the Key Laboratory of Computing Power Network and Information Security, Ministry of Education (No.2023ZD029). Dr. Yang Zhang from the University of North Texas did not receive financial support for this work from the above grants or from any other external project; his contributions were carried out as independent academic research. The authors sincerely thank all colleagues from the participating universities for their valuable collaboration and insights.

References

[1] Z. Gao, T. Shen, Z. Mai, M. R. Bouadjenek, I. Waller, A. Anderson, R. Bodkin, and S. Sanner, "Mitigating the filter bubble while maintaining relevance: Targeted diversification with vae-based recommender systems," in *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '22, p. 2524–2531, 2022.

[2] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua, "Neural collaborative filtering," in *Proceedings of the 26th International Conference on World Wide Web*, WWW '17, (Republic and Canton of Geneva, CHE), p. 173–182, 2017.

[3] Y. Zhang, Y. Wang, K. Wang, Q. Z. Sheng, L. Yao, A. Mahmood, W. E. Zhang, and R. Zhao, "When large language models meet citation: A survey," *arXiv preprint arXiv:2309.09727*, 2023.

[4] L. Wu, P. Sun, Y. Fu, R. Hong, X. Wang, and M. Wang, "A neural influence diffusion model for social recommendation," in *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR'19, (New York, NY, USA), p. 235–244, 2019.

[5] Y. Yang, L. Wu, Y. Liao, Z. He, P. Shao, R. Hong, and M. Wang, "Invariance matters: Empowering social recommendation via graph invariant learning," in *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 2038–2047, 2025.

[6] Y. Yang, L. Wu, Z. Wang, Z. He, R. Hong, and M. Wang, "Graph bottlenecked social recommendation," in *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 3853–3862, 2024.

[7] M. Fernández and A. Bellogín, "Recommender systems and misinformation: The problem or the solution?," in

Proceedings of the Workshop on Online Misinformation- and Harm-Aware Recommender Systems co-located with 14th ACM Conference on Recommender Systems, vol. 2758, pp. 40–50, 2020.

[8] D. Zhang, L. Qi, Y. Liu, X. Zhou, J. Xie, H. Xiang, X. Xu, X. Zhang, Y. Cao, and Y. Zhang, "Hrcformer: Hierarchical recursive convolution-transformer with multi-scale adaptive recalibration for time series forecasting," in *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*, (New York, NY, USA), p. 4076–4085, 2025.

[9] C. R. Sunstein, "Infotopia: How many minds produce knowledge," 2006.

[10] J. Piao, J. Liu, F. Zhang, J. Su, and Y. Li, Humanai adaptive dynamics drives the emergence of information cocoons, *Nat. Mac. Intell.*, vol. 5, no. 11, pp. 1214–1224, 2023.

[11] N. J. Stroud, Polarization and partisan selective exposure, *Journal of Communication*, vol. 60, no. 3, pp. 556–576, 2010.

[12] M. Sülflow, S. Schäfer, and S. Winter, Selective attention in the news feed: An eye-tracking study on the perception and selection of political news posts on facebook, *New Media & Society*, vol. 21, no. 1, pp. 168–190, 2019.

[13] Y. Xi, M. Weng, W. Chen, C. Yi, D. Chen, G. Guo, M. Zhang, J. Wu, Y. Jiang, Q. Liu, Y. Yu, and W. Zhang, "Bursting filter bubble: Enhancing serendipity recommendations with aligned large language models," *CoRR*, vol. abs/2502.13539, 2025.

[14] Y. Xu, E. Wang, Y. Yang, and H. Xiong, GS-RS: A generative approach for alleviating cold start and filter bubbles in recommender systems, *IEEE Trans. Knowl. Data Eng.*, vol. 36, no. 2, pp. 668–681, 2024.

[15] C. Gao, S. Wang, S. Li, J. Chen, X. He, W. Lei, B. Li, Y. Zhang, and P. Jiang, "Cirs: Bursting filter bubbles by counterfactual interactive recommender system," *ACM Trans. Inf. Syst.*, vol. 42, no. 1, 2023.

[16] N. Li, C. Gao, J. Piao, X. Huang, A. Yue, L. Zhou, Q. Liao, and Y. Li, "An exploratory study of information cocoon on short-form video platform," in *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pp. 4178–4182, 2022.

[17] Y. Ge, S. Zhao, H. Zhou, C. Pei, F. Sun, W. Ou, and Y. Zhang, "Understanding echo chambers in ecommerce recommender systems," in *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, pp. 2261–2270, 2020.

[18] T. T. Nguyen, P.-M. Hui, F. M. Harper, L. Terveen, and J. A. Konstan, "Exploring the filter bubble: the effect of using recommender systems on content diversity," in *Proceedings of the 23rd International Conference on World Wide Web*, p. 677–686, 2014.

[19] H. Yuan, Y. Zhang, Q. Z. Sheng, L. Yao, Y. Zhou, X. He, and Z. Wang, "Pkgrec: Personal knowledge graph construction and mining for federated recommendation enhancement," in *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*, (New York, NY, USA), p. 3973–3982, 2025.

[20] C. Aslay, A. Matakos, E. Galbrun, and A. Gionis, "Maximizing the diversity of exposure in a social network," in *2018 IEEE International Conference on Data Mining* (*ICDM*), pp. 863–868, 2018.

[21] Y. Zhang, Y. Wang, Q. Z. Sheng, A. Mahmood, W. Emma Zhang, and R. Zhao, "Tdm-cfc: Towards document-level multi-label citation function classification," in *International Conference on Web Information Systems Engineering*, pp. 363–376, Springer, 2021.

[22] Y. Zhang, Y. Wang, Q. Z. Sheng, A. Mahmood, W. E. Zhang, and R. Zhao, "Hybrid data augmentation for citation function classification," in *2023 International joint conference on neural networks* (*IJCNN*), pp. 1–8, IEEE, 2023.

[23] Y. Zhang, R. Zhao, Y. Wang, H. Chen, A. Mahmood, M. Zaib, W. E. Zhang, and Q. Z. Sheng, Towards employing native information in citation function classification, *Scientometrics*, vol. 127, no. 11, pp. 6557–6577, 2022.

[24] L. C. Freeman, et al, Centrality in social networks: Conceptual clarification, *Social network*: *critical concepts in sociology*. *Londres*: *Routledge*, vol. 1, pp. 238–263, 2002.

[25] S. Burt Ronald, "Structural holes: the social structure of competition," 1992.

[26] Z. Lin, Y. Zhang, Q. Gong, Y. Chen, A. Oksanen, and A. Y. Ding, Structural hole theory in social network analysis: A review, *IEEE Transactions on Computational Social Systems*, vol. 9, no. 3, pp. 724–739, 2022.

[27] L. Breiman, Random forests, *Machine learning*, vol. 45, pp. 5–32, 2001.

[28] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pp. 785–794, 2016.

[29] L. Prokhorenkova, G. Gusev, A. Vorobev, A. V. Dorogush, and A. Gulin, "Catboost: unbiased boosting with categorical features," *Advances in neural information processing systems*, vol. 31, 2018.

[30] J. Zhang, B. Liu, J. Tang, T. Chen, and J. Li, "Social influence locality for modeling retweeting behaviors," in *Twenty-third international joint conference on artificial intelligence*, Citeseer, 2013.

[31] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.

[32] L. Yang, Z. Liu, Y. Wang, C. Wang, Z. Fan, and P. S. Yu, "Large-scale personalized video game recommendation via social-aware contextualized graph neural network," in *Proceedings of the ACM Web Conference*, p. 3376–3386, 2022.

[33] M. O'Neill, E. Vaziripour, J. Wu, and D. Zappala, "Condensing steam: Distilling the diversity of gamer behavior," in *Proceedings of the 2016 Internet Measurement Conference*, (New York, NY, USA), p. 81–95, 2016.

[34] M. Gu, T.-F. Zhao, L. Yang, X.-K. Wu, and W.-N. Chen, Modeling information cocoons in networked populations:

Insights from backgrounds and preferences, *IEEE Transactions on Computational Social Systems*, vol. 11, no. 3, pp. 4497–4510, 2024.

[35] W. Jiang, X. Gao, G. Xu, T. Chen, and H. Yin, “Challenging low homophily in social recommendation,” in *Proceedings of the ACM Web Conference 2024*, p. 3476–3484, 2024.

[36] A. Edwards, (how) do participants in online discussion forums create ‘echo chambers’?: The inclusion and exclusion of dissenting voices in an online forum about climate change, *Journal of Argumentation in Context*, vol. 2, no. 1, pp. 127–150, 2013.

[37] S. Flaxman, S. Goel, and J. M. Rao, Filter bubbles, echo chambers, and online news consumption, *Public opinion quarterly*, vol. 80, no. S1, pp. 298–320, 2016.

[38] M. Cinelli, G. D. F. Morales, A. Galeazzi, W. Quattrociocchi, and M. Starnini, The echo chamber effect on social media, *Proc. Natl. Acad. Sci. USA*, vol. 118, no. 9, p. e2023301118, 2021.

[39] H. Lin, J. Jia, J. Qiu, Y. Zhang, G. Shen, L. Xie, J. Tang, L. Feng, and T.-S. Chua, Detecting stress based on social interactions in social networks, *IEEE Transactions on Knowledge and Data Engineering*, vol. 29, no. 9, pp. 1820–1833, 2017.

[40] H. Wang, Q. Meng, J. Fan, Y. Li, L. Cui, X. Zhao, C. Peng, G. Chen, and X. Du, “Social influence does matter: User action prediction for in-feed advertising,” in *Proceedings of the AAAI Conference on artificial intelligence*, vol. 34, pp. 246–253, 2020.

[41] E. Bakshy, S. Messing, and L. A. Adamic, Exposure to ideologically diverse news and opinion on facebook, *Science*, vol. 348, no. 6239, pp. 1130–1132, 2015.

[42] M. Orbach, Y. Bilu, A. Toledo, D. Lahav, M. Jacovi, R. Aharonov, and N. Slonim, “Out of the echo chamber: Detecting countering debate speeches,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7073–7086, 2020.

[43] A. J. Chaney, B. M. Stewart, and B. E. Engelhardt, “How algorithmic confounding in recommendation systems increases homogeneity and decreases utility,” in *Proceedings of the 12th ACM Conference on Recommender Systems*, pp. 224–232, 2018.

[44] A. Anderson, L. Maystre, I. Anderson, R. Mehrotra, and M. Lalmas, “Algorithmic effects on the diversity of consumption on spotify,” in *The Web Conference*, pp. 2155–2165, 2020.

[45] S. Du and S. Gregory, “The echo chamber effect in twitter: does community polarization increase?,” in *Complex Networks & Their Applications V - Proceedings of the 5th International Workshop on Complex Networks and their Applications*, vol. 693, pp. 373–378, 2016.

[46] H. Peng and C. Liu, “Breaking the information cocoon: When do people actively seek conflicting information?,” in *Information: Equity*, Diversity, Inclusion, Justice, and Relevance - Proceedings of the 84th ASIS&T Annual Meeting, vol. 58, pp. 801–803.

[47] W. Liu and W. Zhou, Research on solving path of negative effect of “information cocoon room” in emergency, *Discrete Dynamics in Nature and Society*, vol. 2022, no. 1, p. 1326579, 2022.

[48] Y. Jeon, J. Kim, S. Park, Y. Ko, S. Ryu, S.-W. Kim, and K. Han, “Hearhere: Mitigating echo chambers in news consumption through an ai-based web system,” *Proc. ACM Hum.-Comput. Interact.*, vol. 8, no. CSCW1, 2024.

Shaokui Gu received the BS degree from the School of Information Science and Engineering, Hebei University of Science and Technology, China, in 2021. He received the MS degree at the School of Software, Dalian University of Technology, China, in 2024. He is currently working toward the PhD degree in Computer Science at the School of Computer Science, Fudan University. His current research interests include data mining, deep learning, and social networks.

Xu Yuan received the B.S. degree in information management from Carnegie Mellon University, United States, in 2000. He is currently an Associate Professor with the Dalian University of Technology, Dalian, China. His research interests are big data visualization analysis, mobile Internet, financial data mining, and electronic commerce.

Yang Zhang received the Ph.D. degree in Computer Science from the School of Computing, Macquarie University, Sydney, NSW, Australia, in 2023. He is a Clinical Assistant Professor in Data Science within the Department of Data Science at the University of North Texas, USA. His research interests lie at the intersection of artificial intelligence and the Artificial Intelligence of Things (AIoT), with a particular focus on natural language processing, recommender systems, and federated learning. Dr. Zhang’s research work explores the application of Large Language Models (LLMs) for advanced data mining and the development of intelligent IoT systems. His research has been published in several recognized venues e.g., AAAI, WWW, ICSE, UBICOMP, INFOCOM, ICSOC, ICWS, CIKM, ICDM, IEEE Transactions on Services Computing, Scientometrics, and Journal of Informetrics.

Xiangguo Sun received the PhD degree from Southeast University, China, in 2022. He is a postdoctoral research fellow with the Chinese University of Hong Kong. He was recognized as the ”Social Computing Rising Star” in 2023 by CAAI. He studied with Zhejiang Lab as a visiting researcher in 2022. In the same year, he won the Distinguished PhD Dissertation Award. During his PhD study, he worked as a research intern with Microsoft Research Asia (MSRA) from 2021 to 2022 and won the ”Award of Excellence”. He studied as a joint PhD student with The University of Queensland hosted by ARC Future Fellow Prof. Hongzhi Yin from 2019 to 2021. His research interests include social computing and network learning. He was the winner of the Best Research Paper Award at KDD’23, which is the first time for Mainland and Hong Kong of China. He has published 12 CORE A*, 10 CCF A, and 15 SCI (including 8 IEEE Trans), some of which appear in SIGKDD, VLDB, The Web Conference (WWW), IEEE Transactions on Knowledge and Data Engineering, ACM Transactions on Information Systems, WSDM, IEEE Transactions on Neural Networks and Learning Systems, CIKM, etc.

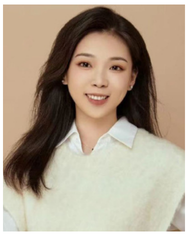

Shuo Yu received the PhD degree from Dalian University of Technology, China, in 2022. She is currently an Associate Professor with the School of Computer Science and Technology, Dalian University of Technology, Dalian, China. She has authored over 50 papers in ACM/IEEE conferences, journals, and magazines. Her research interests include network science, data science, computational social science, and knowledge science. Dr. Yu is a member of the ACM. She is also a member of the IEEE Young Professionals, IEEE Computational Intelligence Society, IEEE Standards Association, and IEEE Women in Engineering. She received several awards, including IEEE DataCom 2017 Best Paper Award, IEEE CSDE 2020 Best Paper Award, and ACM/IEEE JCDL 2020 The Vannevar Bush Best Paper Honorable Mention. She has served as the track chair and a PC/TPC member for four international conferences. She received more than 1000 citations for her work. She has also given keynote speeches and academic reports at nine international conferences, including the International Workshop on Machine Learning on Graphs (MLoG).

Qiang Duan (Senior Member, IEEE) received the Ph.D. in Electrical Engineering from the University of Mississippi, United States, in 2003. He is currently a professor of Information Sciences and Technology with Pennsylvania State University Abington College. His current research interests include networkedge-cloud convergence, cognitive and autonomous networking, and ubiquitous intelligence in the future Internet. He has published more than 100 research papers in these areas and served on the editorial boards of various research journals in the field of networking and distributed computing.

Fenghua Tong received the M.S. degree in mathematics from the Civil Aviation University of China, Tianjin, China, in 2017, and the Ph.D. degree in information security from Beijing University of Posts and Telecommunications, Beijing, China, in 2021. He is currently an Associate Researcher with Shandong Provincial Key Laboratory of Computer Networks, Shandong Computer Science Center (National Supercomputer Center in Jinan), Qilu University of Technology (Shandong Academy of Sciences), Jinan, China. His research interests include compressive sensing and image encryption.

Qingyuan Gong received her PhD degree in Computer Science at School of Computer Science, Fudan University in 2020, and worked as a Postdoc Researcher in the same school until 2022. She is currently a tenure-track Associate Professor at the Research Institute of Intelligent Complex Systems, Fudan University. Her research interests include network security, user behavior analysis and computational social systems. She published referred papers in IEEE Communications Magazine, ACM TWEB, IEEE TDSC, IEEE TKDE, IEEE TMC, Springer WWW Journal, ACM CIKM and ICPP.