

The Relationship Between User Interest Diversity and Information Cocoon Formation Speed in Long-Term Recommendation Environments

Xinyu Zhang

*Joint Institute of Technology, Guangzhou College of Technology and Business, Foshan, China
zxy.817799@gmail.com*

Abstract. Personalized recommender systems may repeatedly expose users to similar content, potentially contributing to information cocoons. Existing research has mainly examined how recommendation algorithms influence content narrowing, while the role of persistent user characteristics remains less explored. This study investigates whether users with different initial levels of interest diversity experience different rates of information cocoon formation in a long-term recommendation environment. Using MovieLens 1M, interest diversity is measured with Shannon entropy over genre distributions. Users are classified into low-, medium-, and high-diversity groups, and 900 users are selected through stratified random sampling. A 20-round UserCF recommendation simulation is then conducted. Mean interest entropy decreases in all three groups, but at different rates. The high-diversity group exhibits the fastest decline, at approximately four times the rate of the low-diversity group. Contrary to the initial expectation, the low-diversity group declines most slowly. This may reflect a floor effect, because its initially concentrated interest distribution leaves less room for further reduction. These findings show that initial interest diversity is associated with the speed of information cocoon formation, shifting attention from whether cocoons form to how quickly they develop.

Keywords: Interest Diversity, Information Cocoon, Information Cocoon Formation, Recommender Systems, User-Based Collaborative Filtering

1. Introduction

Recommender systems are a core technology for content distribution on digital platforms, helping users identify relevant content amid information overload [1, 2]. By analyzing users' historical behavior and preferences, these systems generate personalized recommendations that improve information access and user experience. However, repeatedly prioritizing content similar to users' existing interests may gradually narrow what they encounter. This concern is commonly discussed through the related concepts of information cocoons and filter bubbles, which describe increasingly homogeneous information environments with fewer opportunities to access diverse content and viewpoints [3, 4].

Previous studies have mainly examined this problem from the perspective of recommendation mechanisms, but their findings are not uniform. Recommendation may narrow some forms of content diversity while increasing others, and user choices can produce narrowing even without recommender systems [5-8]. These findings indicate that information cocoon formation is shaped by the interaction between recommendation mechanisms and user behavior rather than by algorithms alone.

Evidence from Spotify suggests that consumption diversity remains relatively stable over time, supporting the view that interest diversity can be considered a persistent user characteristic [9]. Recent work has also modeled information cocoons through long-term user–system simulations, but has mainly compared recommendation algorithms and mitigation methods rather than differences among users [10]. Therefore, this paper asks: In a long-term recommendation environment, do users with different initial levels of interest diversity form information cocoons at different speeds? By focusing on formation speed, this study shifts attention from whether information cocoons form to how quickly they form and examines the role of user-level differences more directly.

2. Literature review and research gap

2.1. Recommender systems and information cocoons

Recommender systems generate personalized recommendations based on users' historical behavior and preferences [1, 2]. This mechanism improves content matching, but reliance on past behavior may narrow users' information exposure. Pariser's filter bubble emphasizes algorithmic filtering that users may not easily detect, which increases exposure to content consistent with existing preferences [3]. Sunstein's information cocoon emphasizes active user choice: individuals may gradually reduce their exposure to diverse information and viewpoints [4]. Both involve narrowing exposure, but the former emphasizes algorithms, while the latter also includes user choice. This paper conceptualizes information cocoon formation as a dynamic process in which users' interest distributions become concentrated during long-term interactions with a recommender system.

2.2. Algorithmic effects and user behavior

Existing studies have not reached a consensus on whether recommender systems necessarily reduce content diversity. Using longitudinal MovieLens data, Nguyen et al. found slight declines in the diversity of recommended and consumed content over time. However, users who consumed recommended items experienced less narrowing than users who did not [5]. Fleder and Hosanagar found that some recommendation mechanisms may reinforce popular items and reduce aggregate sales diversity, yet individual-level diversity may increase as users encounter new products [6].

At the group level, Hosanagar et al. found no evidence that personalization necessarily divides users into isolated consumption groups. Personalization may increase individual consumption diversity while making consumption patterns more similar across users [7]. These findings indicate that diversity at the individual, between-user, and market levels should not be treated as the same phenomenon. Content narrowing is not caused solely by algorithms. Aridor et al. showed that users' own decision-making may narrow their content choices even without a recommender system. Recommendation mechanisms may reduce the natural narrowing of individual users' content choices while increasing homogeneity across users [8]. Thus, information cocoon formation is shaped by both recommendation mechanisms and user behavior; algorithm comparisons alone cannot fully explain how different users change during long-term recommendation interactions.

2.3. Interest diversity and research gap

Among user characteristics, this paper focuses on interest diversity because of its relative stability. Anderson et al. found that Spotify users' consumption diversity remained relatively stable over time [9]. Although the measures differ, this result supports the view that users' initial interest structures are relatively stable.

Zhang et al. compared information cocoon formation and mitigation across algorithms in multi-round simulations [10], but did not examine user differences within the same mechanism. This study therefore compares low-, medium-, and high-diversity users in the same long-term UserCF environment.

3. Methodology

3.1. Dataset

The MovieLens 1M dataset was used. The original *movies.dat* file contained 3,883 movie records, while *ratings.dat* contained 1,000,209 explicit ratings from 6,040 users on a scale from 1 to 5. After an inner join by MovieID, the analytical dataset covered 3,706 movies with ratings. Movie genres were used to construct user interest distributions, while demographic variables such as gender, age, and occupation were excluded.

3.2. Measurement of interest diversity

Interest diversity was measured using Shannon entropy based on genre distributions. For a multi-genre movie, its total contribution was divided equally among all listed genres. Each user's contributions were aggregated into a genre preference distribution; rating values did not affect genre contributions. Genre weights were normalized into probabilities p_i , and interest entropy was calculated using the natural logarithm:

$$H(u) = -\sum_{i=1}^G p_i \ln p_i \quad (1)$$

where G is the number of genres in the user's interest distribution, and p_i is the proportion of genre i for user u . Higher entropy indicates a more dispersed interest distribution.

3.3. User grouping and sampling

All 6,040 users were ranked by interest entropy in ascending order. The lowest 30% were assigned to the low-diversity group (1,812 users), the middle 40% to the medium-diversity group (2,416 users), and the highest 30% to the high-diversity group (1,812 users). Stratified random sampling selected 300 users from each group, producing 900 users in total. A fixed random seed of 42 was used for sampling and acceptance to support reproducibility.

3.4. Recommendation model and long-term simulation

This study used a UserCF-based recommendation simulation with interest-based score adjustment. A user-item matrix was constructed from the original ratings, and cosine similarity was used to calculate user similarity. Candidate movies came from the 100 most similar neighbors. Movies

already rated by the target user in the original dataset were excluded, and only items receiving neighbor ratings of at least 4 were retained.

The Interest Score was the sum of the user's current interest weights for all genres of a candidate movie:

$$I(u, m) = \sum_{g \in G_m} w_{u,g} \quad (2)$$

The final recommendation score was:

$$S(u, m) = CF(u, m) + 20I(u, m) \quad (3)$$

where $CF(u, m)$ is the sum of similarity-weighted neighbor ratings for movie m , G_m is its genre set, and $w_{u,g}$ is user u 's current interest weight for genre g . In each round, the top 10 movies by final score were recommended.

The simulation lasted 20 rounds. Each recommended movie was independently accepted with probability 0.5, regardless of its score. Each accepted movie contributed equally across its genres to the user's interest distribution at a learning rate of 0.1, after which the distribution was renormalized. Accepted movies were not written back to the user-item matrix. The matrix and user similarity therefore remained fixed, while only the current interest distribution changed across rounds.

3.5. Formation speed and statistical analysis

Interest entropy was recorded from Round 0, before recommendation, to Round 20 for every user, producing 18,900 records. The net decline in interest entropy was used as the operational indicator of information cocoon formation. Formation speed was defined as:

$$FS(u) = \frac{H_{u,0} - H_{u,20}}{20} \quad (4)$$

where $H_{u,0}$ and $H_{u,20}$ are the entropy values at Rounds 0 and 20. Larger positive values indicate higher formation speed, while negative values indicate a net increase in interest diversity after 20 rounds.

One-way ANOVA tested overall differences in formation speed across the three diversity groups, followed by Tukey HSD pairwise comparisons ($\alpha = 0.05$).

4. Results

4.1. Changes in interest entropy across recommendation rounds

Figure 1 shows mean interest entropy from Round 0 to Round 20 for the three diversity groups. Mean entropy declined in all three groups, with the largest decrease in the high-diversity group and the smallest in the low-diversity group (Table 1).

The medium and high groups showed an overall steady decline. In the low group, mean entropy rose at Round 1 and first fell below its initial level at Round 4, after which it declined overall.

At Round 20, the low-to-medium-to-high ordering remained, but the gaps between groups had narrowed.




Figure 1. Mean interest entropy across recommendation rounds

4.2. Formation speed across diversity groups

Table 1 and Figure 2 show that mean formation speed increased across the low-, medium-, and high-diversity groups. The mean for the high-diversity group was approximately four times that for the low-diversity group.

Table 1. Entropy changes and formation speed across diversity groups

Group	Mean Entropy at Round 0	Mean Entropy at Round 20	Entropy Decline	Mean Formation Speed	Formation Speed SD
Low Diversity	1.8085	1.6865	0.1219	0.0061	0.0149
Medium Diversity	2.1670	1.8689	0.2980	0.0149	0.0117
High Diversity	2.3790	1.8986	0.4804	0.0240	0.0130

Note: n = 300 per group. SD = standard deviation of formation speed.




Figure 2. Mean formation speed across diversity groups

Note: Error bars show SEM; n = 300 per group.

Negative formation-speed values were observed for 122 of 900 users (13.6%): 100 low-, 19 medium-, and 3 high-diversity users. Thus, individual trajectories differed despite positive group means.

4.3. Statistical differences between groups

A one-way ANOVA tested formation speed differences across the three groups. The overall difference was statistically significant, $F(2, 897) = 137.08$, $p < .001$.

All Tukey HSD pairwise comparisons were significant ($p < .001$): high exceeded low by 0.0179 and medium by 0.0091, while medium exceeded low by 0.0088.

5. Discussion

5.1. Interpretation of the main findings

The results showed that information cocoon formation speed differed across initial interest diversity levels, increasing across the low-, medium-, and high-diversity groups. This was contrary to the earlier expectation that low-diversity users would form information cocoons faster, but the association does not establish a direct effect.

The high group's faster formation speed may reflect its more dispersed initial interests, which left greater room for entropy to decline, while the low group's lower initial entropy left less room for further decline. Repeated recommendations and interest updating may also have concentrated interests. Because acceptance was random, some low-diversity users may have accepted movies outside their main genres, temporarily increasing entropy. The narrower Round 20 gaps indicated partial convergence.

5.2. Comparison with previous findings and implications

The increase in interest entropy among some low-diversity users is consistent with [5]: recommendation exposure does not necessarily narrow user diversity. These individual-level findings do not address market-level diversity in [6] or between-user fragmentation in [7].

The decline in all groups contrasts with the within-user alleviation reported in [8], possibly because of different measures and contexts. The association between initial characteristics and long-term trajectories echoes [9], while the group comparison extends the long-term simulation in [10].

The same recommendation mechanism produced different long-term outcomes across users, so overall averages may hide user heterogeneity. High initial diversity should therefore not automatically be treated as low risk.

5.3. Limitations

Behavioral realism is limited by the fixed 0.5 acceptance probability, static user-item matrix and similarity, omitted feedback from accepted movies, possible repeated recommendations, and fixed parameters. Genre-based entropy is a coarse proxy, one dataset limits generalizability, ANOVA assumptions were not formally tested, and no baseline was included.

6. Conclusion

This study examined the relationship between initial interest diversity and information cocoon formation speed in a long-term recommendation environment. Using MovieLens 1M, genre-based Shannon entropy measured interest diversity. Nine hundred users were divided into low-, medium-, and high-diversity groups for a 20-round recommendation simulation. Formation speed was calculated from the net decline in interest entropy.

Mean interest entropy ended below its initial level in all three groups, although the declines differed. Formation speed also differed significantly across groups, $p < .001$, increasing from low to high. The high group mean was approximately four times the Low group mean. This was contrary to the initial expectation that users with lower interest diversity would form information cocoons faster. Higher initial interest entropy allowed more room for decline, which may partly explain the high group's faster formation speed. Some low-diversity users instead showed increased interest entropy, indicating varied within-group trajectories. At Round 20, the original ordering remained, but group gaps narrowed. The findings indicate an association rather than a confirmed direct causal relationship. The same recommendation mechanism may produce different long-term outcomes across users. Overall averages may hide user heterogeneity, while formation speed can complement endpoint-based analyses of interest diversity. It may therefore be inappropriate to regard high-diversity users as low-risk solely because their initial interests are more diverse.

The study is limited by a fixed acceptance probability, static user similarity, possible repeated recommendations, fixed parameters, and one genre-based measure. Future research could use more realistic acceptance models, dynamically updated data, sensitivity analyses of key parameters, and comparisons across algorithms and datasets.

References

- [1] Ricci, F., Rokach, L., & Shapira, B. (2022). *Recommender systems handbook* (3rd ed.). Springer.
<https://doi.org/10.1007/978-1-0716-2197-4>
- [2] Zhang, S., Yao, L., Sun, A., & Tay, Y. (2019). Deep learning-based recommender system: A survey and new perspectives. *ACM Computing Surveys*, 52(1), Article 5, 38.
- [3] Pariser, E. (2011). *The filter bubble: What the internet is hiding from you*. Penguin Press.
- [4] Sunstein, C. R. (2007). *Republic.com 2.0*. Princeton University Press.
<https://press.princeton.edu/books/hardcover/9780691133560/republiccom-20>
- [5] Nguyen, T. T., Hui, P.-M., Harper, F. M., Terveen, L., & Konstan, J. A. (2014). Exploring the filter bubble: The effect of using recommender systems on content diversity. In *Proceedings of the 23rd International Conference on World Wide Web (WWW'14)*. ACM, 677-686.
- [6] Fleder, D., & Hosanagar, K. (2009). Blockbuster culture's next rise or fall? The impact of recommender systems on sales diversity. *Management Science*, 55(5), 697-712.
- [7] Hosanagar, K., Fleder, D., Lee, D., & Buja, A. (2014). Will the global village fracture into tribes? Recommender systems and their effects on consumer fragmentation. *Management Science*, 60(4), 805-823.
- [8] Aridor, G., Gonçalves, D., & Sikdar, S. (2020). Deconstructing the filter bubble: User decision-making and recommender systems. In *Proceedings of the 14th ACM Conference on Recommender Systems (RecSys'20)*. ACM, pp.82-91.
- [9] Anderson, A., Maystre, L., Anderson, I., Mehrotra, R., & Lalmas, M. (2020). Algorithmic effects on the diversity of consumption on Spotify. In *Proceedings of The Web Conference 2020 (WWW'20)*. ACM, pp.2155-2165.
- [10] Zhang, L., Lian, Y., Wu, H., Song, C., & Yuan, X. (2025). An exploratory study on information cocoon in recommender systems. *Data Science and Engineering*, 10(3), 495-514.